



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 971 519 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
12.01.2000 Bulletin 2000/02

(51) Int Cl. 7. H04L 29/14, H04L 1/18

(21) Application number: 99304516.0

(22) Date of filing: 09.06.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 06.07.1998 US 110218

(71) Applicant: Hewlett-Packard Company
Palo Alto, California 94304 (US)

(72) Inventors:
• Miller, Christopher S.
Ft. Collins, CO 80525 (US)
• Alley, Jeffrey D.
Ft. Collins, CO 80526 (US)
• Titze, Thomas
71154 Nufringen (DE)

(74) Representative: Powell, Stephen David et al
WILLIAMS, POWELL & ASSOCIATES
4 St Paul's Churchyard
London EC4M 8AY (GB)

(54) Computer data transfer

(57) A transport manager (106, 112, 1204, 1254) manages the transport of data between two secured systems (100, 102, 1200, 1250) via various transport media and techniques (1300..1354). The transport media include media and methods which are not precluded from data exchange by security measures implemented on the source and destination systems, e.g. e-mail network communication (1300), file transfer protocol network communication (1302), hypertext transfer protocol (1304), and manual exchange via magnetic or optical

media (1308). Each desired exchange identifies a particular source and destination computing system (100, 104, 102, 114, 1200, 1250, 1252). The methods for transmitting and receiving the data identify the requested exchange as a job (1400..1404), associate a job ID with the job (1503), and track the status of the job in a job database file (1406..1408, 1508, 1600..1622). An address book table is used to look up preferred transport media, methods, and parameters associated with a destination system (1504).




FIG. 1

Description

[0001] The present invention relates to computer network communications and in particular to methods and structures for simple yet robust and reliable file exchange between otherwise secured systems using a variety of communication media.

[0002] Computer networks are used for interconnection of a plurality of computer systems for purposes of exchanging information between such systems. In general, a network includes a number of computer systems interconnected by a one or more communication links. Local area network (LAN) is a term generally applied to relatively smaller networks wherein a number of computer systems are interconnected via a relatively short distance high-speed communication links. For example, a LAN may interconnect a number of computers within an office, a building or a campus. By contrast, wide area networks (WANs) tend to be larger networks comprising a larger number of computers interconnected by a wide variety of high-speed (shorter distance) and lower speed (longer distance) communication links. For example, a global business enterprise may interconnect their installations around the world with a WAN while each particular installation may utilize a LAN to interconnect its particular computer systems.

[0003] In a local area network it is common that all computers cooperate using the related low-level protocols. For example, all computers on a particular local area network (LAN) would typically share use of a single protocol such as TCP/IP, NETBEUI (from Microsoft), or IPX/SPX (from Novell). It is also common for multiple such protocols to be simultaneously operable over a particular LAN or WAN. Higher level networking protocols and applications then exchange data via these lower-level communication protocols.

[0004] Often a group of computer systems on a network are referred to as an enterprise. In general, as used herein, an enterprise is a large and diverse network connecting most major points in a company or other organization. As used herein, an enterprise may be synonymously referred to as a computing enterprise or a business enterprise.

[0005] It is common in computer networking to view the interconnected communications as a layered model, specifically a seven layer model often referred to as the Open Systems Interconnect (the so called OSI standard maintained by the International Standards Organization (ISO)). Higher layers in the model transfer their information to a corresponding next lower layer for detailed management of lower-level transmission issues. The higher layers of the model therefore communicate in effect with corresponding higher level components in other computers on the network. Each lower level of the OSI model, in turn, communicates with its corresponding lower layer embodied within other computers on the network. Such a layered model allows computer networking software to be developed in accordance with standard programming interfaces between the various layers. A particular layer may therefore be replaced by an equivalent component which provides similar services to the layer above and invokes lower-level services from a lower layer in accordance with a standard interface definition.

[0006] In general, wide area networks (WANs) use a wider variety of physical communication media for the exchange of data among a more geographically disperse set of computing systems. Each such diverse lower-level physical communication medium utilizes a corresponding lower level protocol layer in, for example, the OSI model.

[0007] Security is a constant challenge for networks and computing engineers responsible for networks. In particular, in wide area network applications, it is important for computing systems attached to such a network to secure their resources from inappropriate, unauthorized access. The Internet is an example of a global wide area network where security measures are often critical to an ongoing business enterprise connected to the Internet. Such security measures are required to assure that unauthorized third parties, anywhere in the world, cannot gain access to sensitive materials within the enterprise via the global, publicly accessible, Internet.

[0008] Though such security measures (often referred to as firewalls) are vital to secure each particular enterprise, they're very existence creates the burden for those trying to legitimately exchange information between enterprises via such global, public networks. A user in one particular computing enterprise encounters a number of difficulties exchanging data with another user in a different computing enterprise via computer system to computer system network communication links. Though the communication capability may exist, for example via the Internet, safeguards and security measures (firewalls) within each enterprise makes such enterprise to enterprise exchanges difficult - exactly as they are intended to do.

[0009] In general such firewall security measures operate at lower layers of the network communication layered model to filter out potentially harmful network data exchange. For example, the firewall may permit certain protocols to be exchanged only among certain network devices known to be physically secured within the enterprise. Network devices not within the permitted scope of secured devices are not permitted to use the filtered protocols. Should such un-authorized devices attempt such communications, the firewall simply discards their network data transfer requests.

[0010] A significant number of standard networking applications are capable of exchanging certain types of information without interference from security and safeguard mechanisms within an enterprise. Such applications are generally considered secure in that they cannot, in general, cause harm to or disrupt operation of the receiving computer system. For example, electronic mail messages are generally considered benign and therefore readily exchanged between most systems on the Internet regardless of implemented security measures and safeguards within each enterprise.

Other standard protocols including hypertext transfer protocol (HTTP) are similarly limited in their capability and may therefore allow further exchange of information between computing systems which are otherwise secured. Still other "transport" techniques may be utilized to exchange information among otherwise secured business enterprises. For example, manual procedures such as transferring information via magnetic or optical medium, or other direct communication links such as modems which bypass secured network connections may be used to exchange information among the otherwise secured computing enterprises.

[0011] It is generally known to use such widely available networking protocols for purposes of exchanging information between otherwise secured computing enterprises. For example, attachment files are often transmitted with email messages to permit additional information to be transmitted with the known message. Such attachment files may, in general, contain any further information (textual or raw binary data) relating to the underlying email message.

[0012] However, use of such higher level protocols (email exchange protocols) or manual transmission techniques (i.e., exchange via a magnetic or optical medium) does not provide the robust reliable exchange of information normally required for exchange of large quantities of information between two otherwise secured computing enterprises. For example, and when using the email messages to exchange large volumes of information via attached files, there are no automated procedures known for assuring that the entire set of attached files are received once, only once, and reassembled in the correct order.

[0013] In view of the above discussion, it is evident that a need exists for improved protocol management to permit the exchange of data between processes on otherwise secured systems. Such systems may be geographically dispersed and interconnected via a wide area networking or generally devoid of network communication interconnection. Such protocol management would ideally assure robust, reliable transport of information between such otherwise secured computing enterprises.

[0014] The present invention solves the above and other problems, thereby advancing the state of useful arts, by providing management methods and associated structures for managing the transport of messages between processes otherwise secured computing enterprises. Further, the present invention performs such transport utilizing any of several transport mechanisms that may be available for transport between the source and destination systems. In particular, the present invention provides protocols and associated data structures to oversee information exchange via assorted methods including networking protocols as well as manual technologies. The methods and structures of the present invention assure robust, reliable transport of such information using a variety of transport protocols and mechanisms without interference from security mechanisms associated with the communicating computing enterprises.

[0015] In particular, the present invention comprises a transport manager which wraps transported messages in a management layer protocol to assure robust reliable transportation of messages via a variety of transport mechanisms. The transport manager in particular provides acknowledgment and re-transmission functions to assure transmitted information is received and combined as necessary in appropriate sequential order and processed correctly by the receiving application.

[0016] Still more specifically the transport manager of the present invention defines each file transfer as a job and maintains information regarding the status of each job in a job database file. The job database file stores information necessary to restart transmission of the particular job, to send acknowledgment relating to a particular job, and to monitor for various timeout parameters necessary to assure reliable transportation.

[0017] A further aspect of the present invention includes an address book which contains information to identify preferred transport methods associated with particular destination addresses. Other information associated with the address book identifies appropriate timeouts for particular transport methods and paths associated with particular destination addresses. The use of the address book provides a "shortcut" for sending processes to provide protocol and configuration parameters to the management portions of the present invention. An address book entry identifies a particular destination by an ID (e.g., a name) and supplies all other parameters for sending data through the management layer of the present invention to a receiving application in another secured system.

[0018] The transmission related module of the present invention resides in and is operable on a computing system desiring to transmit a message to receiving system. In like manner, a receiving related module of the present invention is cooperatively resident on the destination computing system. The modules on each side of the transmission link of the present invention present an application program interface (API) to client application programs. An application engineer therefore uses the transmission and reception API in integrating application programs on each side of the transmission link of the present invention.

[0019] Each process of the transmit/receive pair maintains information about the job in process in a corresponding database associated with that process (i.e., in a transmit job database and in a receive job database). The transmitting and receiving programs (processes) each register information in the job database of the present invention to identify each process engaged in a data exchange session. A transaction ID provides a globally unique identifier (GUID) for storage and retrieval of data in the job databases.

[0020] The transmission job database is also used to store information required for error recovery to assure reliable transmission of the exchanged data. Though the underlying protocols used to bypass firewall security may not provide

the desired reliability the management functions of the present invention serve to assure such reliable transmissions through timeout functions and acknowledgment handshaking. The management layers of the present invention shield the calling application from the details of using a particular protocol and communication medium. Rather, the sending and receiving applications invoke a common API interface regardless of the particular operating system, communication medium, and protocols selected for transmission of the supplied data.

[0021] The job database in the transmission and reception manager portions provide persistent storage of information to permit failure recovery even through system or power failures. When the system is restored to normal operation, the persistent stored data in the database (transmission job database and/or receiving job database) is used to resume operations exchanging data between the secured systems.

[0022] The present invention therefore permits computing systems segregated by security measures (firewalls) to communicate for purposes of exchanging data in a simple yet reliable and robust manner. The present invention also provides a common API interface for the sending and receiving applications regardless of the particular communication medium and protocol selected for transmission of a particular job. In addition, the present invention provides these features in an architecture which is highly portable between a number of commercial computing environments including Microsoft Windows, HPUX and a number of Unix dialects.

[0023] The above, and other features, aspects and advantages of the present invention will become apparent from the following descriptions and attached drawings.

[0024] Figure 1 is a blocks diagram of processes communicating between secured systems using electronic communication media in accordance with the present invention.

[0025] Figure 2 is a blocks diagram of processes communicating between secured systems using manual transport of storage media in accordance with the present invention.

[0026] Figure 3 is a block diagram depicting data flow between processes of the present invention.

[0027] Figure 4 is a block diagram depicting additional detail of data flow of figure 3 within a sending application.

[0028] Figure 5 is a block diagram depicting additional detail of data flow of figure 3 within a sending application.

[0029] Figure 6 is a block diagram depicting additional detail of data flow of figure 3 within the transmission manager.

[0030] Figure 7 is a block diagram depicting additional detail of data flow of figure 3 within the transmission manager.

[0031] Figure 8 is a block diagram depicting additional detail of data flow of figure 3 within the reception manager.

[0032] Figure 9 is a block diagram depicting additional detail of data flow of figure 3 within the reception manager.

[0033] Figure 10 is a block diagram depicting additional detail of data flow of figure 3 within the reception manager.

[0034] Figure 11 is a block diagram depicting additional detail of data flow of figure 3 within the receiving process.

[0035] Figure 12 is a block diagram describing structure of the elements and processes of the present invention to provide communication between otherwise secured systems.

[0036] Figure 13 is a block diagram describing various exemplary protocols and associated transmission media utilized by the present invention.

[0037] Figure 14 is a flowchart describing operation of methods of the present invention operable within the sending application process.

[0038] Figure 15 is a flowchart describing operation of methods of the present invention operable within the transmission manager.

[0039] Figure 16 is a flowchart describing operation of methods of the present invention operable within the transmission manager.

[0040] Figure 17 is a flowchart describing operation of methods of the present invention operable within the reception manager.

[0041] Figure 18 is a flowchart describing operation of methods of the present invention operable within the reception manager and the receiving application.

[0042] While the invention is susceptible to various modifications and alternative forms, a specific embodiment thereof has been shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that it is not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

[0043] Figure 1 is a block diagram depicting secured computing systems 100 and 102 operable in accordance with methods and structures of the present invention. In particular, transmission manager 106 and corresponding reception manager 112, operable within first computing system 100 and second computing system 102, respectively, manage the exchange of information so as to cooperate with firewalls security measures implemented in the secured computing systems 100 and 102. As shown in figure 1, secured computing systems 100 and 102 communicate via network 150.

[0044] Further as shown in figure 1, network 150 is intended to be representative of any network communication medium including a local area networks (LANs) or wide area networks (WANs).

[0044] In particular, sending application 104 operable within first secured computing system 100 is the originator of information to be transferred to receiving application 114 operable within second secured computing system 102. Send-

ing application 104 works in cooperation with transmission manager 106 within first secured computing system 100 to effectuate the robust, reliable exchange of desired information Reception manager 112 operates in conjunction with receiving application 114 to further assure robust, reliable transmission of the desired data between sending application 104 receiving application 114

5 [0045] Transmission manager 106 is preferably implemented as a client process cooperating with reception manager 112 preferably implemented as a server process within their respective computing systems. As is known in the art, such client and server processes may be comprised of a number of sub components such as daemon programs and associated data structures to effectuate the desired client/server exchange.

10 [0046] In general, firewall security measures are configurable to allow and disallow particular types of data exchanges between particular computing systems. For example, first secured computing system 100 may be one of several computers within a first business computing enterprise. Second secured computing system 102 may likewise be one of many computing systems within a second business computing enterprise. Each business computing enterprise would implement firewall security measures to prevent unauthorized outsiders from accessing sensitive information within their respective business computing enterprises. However, there are often legitimate requirements for transmitting 15 information between two such secured business computing enterprises. For example, a user in a first business computing enterprise may be a supplier of computing products while a user in the second computing enterprise may be a consumer of the products. The user in the first business enterprise may need to transfer large volumes of data to the user in the second enterprise such as upgraded or repaired programs or data files.

20 [0047] Presently known solutions do not provide for robust, reliable transmission of such data. For example, email programs and protocols, though capable of exchanging information through firewall security techniques and capable of attaching data files thereto, do not assure robust, reliable application to application delivery and successful processing of data files as may be required in various applications. Rather, the sending user or application transmits the desired data to the receiving user or application but has no assurance that the transmission was successfully processed at its intended destination (the receiving application).

25 [0048] Transmission manager 106 and reception manager 112 of figure 1 therefore provide robust reliable exchange of information between secured computing systems. In particular, as discussed further herein below, transmission manager 106 and reception manager 112 are operable in conjunction with standard network protocols and programs such as email protocols and HTTP protocols to provide the requisite robust, reliable exchange of information through firewalls. In addition, as discussed further below, the transmission manager 106 and reception manager 112 may 30 communicate via other communication media and protocols including manual transmissions such as the manual exchange of magnetic or optical storage media. Essentially, the cooperating transmission manager 106 and reception manager 112 provide a management layer, a wrapper, in which data is exchanged between secured computing systems. This exchange may be effectuated via any of several communication media and protocols. This management layer provides end to end reliable, robust exchange and processing of information between a sending application on 35 the first secured computing system and a receiving application on the second secured computing system.

40 [0049] As shown in figure 1, transmission manager 106 uses network communication protocols to transmit information generated by sending application 104 through firewall 108 directed to second computing system 102 via network 150. The transmission passes through firewall protection 110 associated with secna system 102 and onto reception manager 112. As shown in figure 1, firewall 108 is implemented as software modules operable within first secured computing system 100. Alternatively, firewall 110 of figure 1 is shown as a stand-alone device connected to network 45 150 to provide firewall security measures on behalf of second computing system 102 (and potentially other connected computing systems). Networking devices such as routers and bridges often provide such firewall protection capabilities within their networking functionality. Firewalls 108 and 110 are therefore intended as exemplary of a large variety of well-known firewall techniques and devices. Transmission manager 106 and reception manager 112 are therefore operable to utilize communication media and protocols which are capable of passing through such firewall protection.

50 [0050] As noted above, firewall protection is known in the art to be applicable to a number of communication techniques and media. Figure 2 depicts an alternative embodiment of the present invention wherein magnetic media 200 (e.g., floppy disk, optical storage medium, tape medium, etc.) are utilized to manually transport requested information between computing system 100 and 102. Transmission manager 106 and reception manager 112 are therefore adapted 55 to utilize a wide variety of communication media and protocols for changing information between the secured computing systems 100 and 102. In general, as shown in figure 2, information generated by sending application is written directly to magnetic media 200 by transmission manager 106 without requiring any interactions with firewall protection 108. Likewise, reception manager 112 may write returned information on magnetic media 200 without interactions with firewall 110. In both cases, reading returned information transmitted information from magnetic media 200 preferably passes through firewall protection techniques so as to assure secured communication between secured computing systems 100 and 102.

5 [0051] As can be seen from the above discussion, "firewall" protection techniques, as used herein, refers to a wide variety of security measures and devices useful with secured computing systems to enable and disable exchange of

particular forms of data between particular systems. Transmission manager 106 reception manager 112 are therefore operable to exchange information on behalf of sending application 104 in receiving application 114 via a variety of communication media and protocols through a variety of firewall techniques or devices.

[0052] Figure 12 is a block diagram depicting apparatus and structure of the present invention to provide robust reliable communication between secured computing systems. First secured computing system 1200 communicates with second secured computing system 1250 through firewalls 1216 and 1264, respectively, via communication medium 1290. In particular, sending process 1202 operable within first secured computing system 1200 communicates with corresponding receiving process 1252 operable within second secured computing system 1250. Transmission manager 1204 operable within first secured computing system 1200 manages the transmission of information sent by sending process 1202. The corresponding reception manager 1254, operable within second secured computing system 1250, further manages the exchange of information between the secured computing systems on behalf of sending process 1202 and receiving process 1252.

[0053] Transmission manager 1204 preferably includes acknowledgment processing element 1206 operable to request acknowledgment from reception manager 1254. Acknowledgment processing element 1206 is also operable to monitor reception of the requested acknowledgment message from the reception manager 1254. As noted further herein below, aspects of the protocols exchanged between the secured computing systems via medium 1290 are symmetric. For example, acknowledgment processing element 1206 preferably further includes processing components for responding to acknowledgment requests received from reception manager 1254.

[0054] Transmission manager 1204 preferably further includes failure monitoring element 1208 which determines whether a transmission from transmission manager 1204 to reception manager 1254 was successful or failed. In particular, failure monitoring element 1208 preferably includes a timeout detector element 1209 for detecting timeout conditions while awaiting acknowledgment of a transmitted job. In general, as discussed in further detail herein below, the successful transmission is detected by reception of an appropriate acknowledgment by acknowledgment processing element 1206. A failed transmission is detected by failure monitoring element 1208 detecting a timeout condition through timeout detector element 1209. Such a failure detection triggers a retransmission attempt within transmission manager 1204. Eventual acknowledgment of a transmitted or re-transmitted job is recognized as a successful transmission of the job from first computing system 1200 to second secured computing system 1250. Re-transmitter element 1210 is responsible for performing requisite retransmission of a job in response to detection of a failure by failure monitoring element 1208.

[0055] Transmission manager 1204 stores persistent information regarding status of transmission jobs in transmit job database 1212. Entries 1214 in transmit job database 1212 retain information regarding status and progress of a particular job to be sent to second secured computing system 1250 on behalf of sending process 1202. Such information preferably includes identification of the destination process (i.e., receiving process 1252 in second secured computing system 1250), timeout parameters associated with the transmission medium, the actual job data, and various other parameters required for transmission of the job between the two secured computing systems. Persistent storage of such job information permits failure detection and recovery of the job to enable robust reliable communication between the secured computing systems.

[0056] Reception manager 1254 operable within second computing system 1250 preferably includes acknowledgment processing element 1256 analogous to acknowledgment processing element 1206 within transmission manager 1204 of first secured computing system 1200. In general, acknowledgment processing element 1256 is operable to respond to acknowledgment requests generated by acknowledgment processing element 1206 within transmission manager 1204. As noted herein above, acknowledgment request and response aspects of the protocol between the two computing systems are symmetric under certain conditions. In other words, acknowledgment processing element 1256 is also operable to generate acknowledgment requests directed back toward acknowledgment processing element 1206.

[0057] Reception manager 1254 maintains information in a received job database 1260 to permit failure detection and recovery for communication between the two computing systems. Entries 1262 in received job database 1260 persistently store parameters required by reception manager 1254 to monitor progress and processing of a job transmitted from transmission manager 1204. For example, entries 1262 preferably include information identifying the originating process (i.e., sending process 1202 and transmission manager 1204), the receiving process 1252 identified by sending process 1202 as associated with the transmitted job, and OUTBOX directory structure identification for storage of transmitted job data, and various other job parameters required for robust, reliable change of information between the two secured computing systems.

[0058] As noted above, communication medium 1290 may be any of several communication media and protocols capable of communicating through firewalls 1216 and 1264. Figure 13 depicts a number of such communication media and protocols operable within transmission manager 1204 and reception manager 1254 of figure 12 for communicating via medium 1290. In particular, email protocols 1300 (such as SMTP, POP, Microsoft Exchange, MAPI, etc.) of figure 13 may communicate through firewall protection techniques via a network communication medium 1350. Similarly,

HTTP protocols 1304 may be used over network communication media 1350. Direct modem connections protocols 1306 may be applied to modem media 1352 (i.e., direct modem connections via telephone or other hard-wired communication links). Furthermore, magnetic or optical storage medium and associated read/write devices 1308 may be used to manually transport job data files between two secured computing systems. Such manual transport is depicted in figure 13 as dashed line 1354.

[0059] FTP protocols 1302 of figure 13 are somewhat unique in that their application to secured systems is not standardized. FTP protocols include their own security measures. Particular enterprise installations may permit FTP protocols to pass through firewall protected systems because the FTP security measure are effective to secure sensitive data. Other installations may simply reject FTP protocols at the firewall security as too large a security risk to their enterprise. Those skilled in the art will therefore recognize that FTP protocols may be applied in conjunction with the management features of the present invention. However, application of FTP protocols in conjunction with the present invention may vary depending upon installation parameters of the particular computing enterprise. FTP protocols are easily used in conjunction with the present invention to exchange information within a single secured system. Though the features of the present invention are most useful in communications between secured systems, they may be similarly applied to non-secured communication applications such as within a single secured enterprise or between unsecured systems.

[0060] Figures 1, 2, and 12 show only a transmission manager portion of the invention resident in the first secured computing system and only a reception manager portion in the second secured computing system. Preferably, the two components are in fact implemented as a single API module such that both transmission and reception functions are operable within each system. More specifically, both transmission and reception API functions are invoked in processing within both secured computing systems for purposes of acknowledgment handshakes discussed below. As discussed further below, the transmission manager receives acknowledgment messages from the second system and the reception manager transmits acknowledgment messages to the first.

25 DATA FLOWS IN THE PREFERRED EMBODIMENTS

[0061] Figure 3 describes the transport management of the present invention in a high-level data flow diagram. A sending application 300 creates a job object utilizing API functions of the transmission manager 302 of the present invention. The transmission manager 302 transmits the data associated with the job object over communication medium 304. As noted above, the communication medium and protocols utilized are capable of passing through firewall protection techniques applied to both the transmitting receiving side of the transport manager. The job object data is received by reception manager 306 of the present invention. Transmission manager 302 and reception manager 306 coordinate acknowledgment and error recovery for transmissions there between.

[0062] The job object received by a reception manager 306 is then forwarded to an identified receiving application 308. Acknowledgment and error messages are communicated in the reverse direction: from the receiving application 308 to reception manager 306, thence from reception manager 306 to transmission manager 302, and finally, on demand, status is reported from transmission manager 302 to sending application 300.

[0063] Communication paths 350 and 352 between sending application 300 and transmission manager 302 are preferably implemented as interprocess communication (IPC) links. Communication paths 354 and 356 between transmission manager 302 and reception manager 306 utilize communication medium 304 as selected for a particular job transmission. Communication paths 358 and 360 between reception manager 306 and receiving application 308 are preferably IPC links.

[0064] Figures 4 and 5 provide additional detail of the data flow performed within the sending application 300 when used in conjunction with the transport manager of the present invention. Specifically, figure 4 describes data flow pertaining to operation of sending application 300 to create a job for transmission to a second secured computing system. Sending application 300 first invokes an API function 400 to create an application registration object 402. This object 402 serves to register the sending application with the transmission manager for further interactions. The registration object 402 so created is then used in invocation of an API function 404 to create a courier object 404. Additionally, sending application invokes another API function 406 to create address and policy objects 408 and 410 identifying the intended destination for the job to be transmitted, and other parameters associated with the transmission of the job. The address object 408 and policy object 410 so created are then used in invocation of API function 412 to create a job object 416 describing the transmission job. Subsequent API function calls 414 add particular data files to the transmission job represented by job object 416. A single transmission job may therefore comprise one or more data files to be transmitted to the second secured computing system. The courier object 418 is associated with the particular job object 416 to be transmitted in this job instance. The courier object 418 therefore may be utilized for transmission of a plurality of discrete jobs, each defined and created by the user as described above with respect to API functions 406, 412, and 414. Each particular job object 416 so created is then transmitted by the sending application by invocation of the sending method 420 of the courier object 418. Should the transmission result in error conditions

(an unrecoverable error), an error object 422 is returned to the sending application from the transmission manager.

[0065] Figures 5 provides additional data flow details regarding operation of the sending application for querying the status of a previously transmitted job. API functions 400 and 404 are used as described above for creation of registration object 402 and courier object 418 for purposes of querying the status of a previously transmitted job. Sending application invokes API function 500 to obtain the transaction ID for the previous job for which status is to be queried. The sending application then invokes an API function 502 to create a transaction ID object 504. The sending application then invokes the query method 506 of the courier object 418 to determine the present status of the job identified by transaction ID object 504. The present status of the identified job is returned to the sending application by invocation of the query method 506. When the status of a particular job changes, a new status object 506 is asynchronously received. The status object 508 is then used to update the status of its identified job. By repetitive invocation of the query method 506, the sending application may eventually determine the completion status of a previously transmitted job.

[0066] Figures 6 and 7 provide additional data flow details regarding the operation of transmission manager 302 of figure 3. A courier communication buffer is received from the sending application by element 600. The courier communication buffer includes an input buffer object 602. Preferably, the inter-process communications between the sending application and transmission manager are performed utilizing well-known inter-process socket communication techniques. Those skilled in the art will readily recognize a variety of other design choices for inter-process communication.

[0067] The received input buffer object 602 associated with the courier communication buffer is parsed by operation of element 604 and an appropriate job object command 606 is created by invocation of an appropriate API function. Elements 608 then launches an independently operating thread to process the requested transmission job object command 606.

[0068] Figure 7 provides additional details regarding the data flow within transmission manager 302 to process of a particular job (element 606 of figure 6). Specifically, elements 700 through 711 relate to processing of a new transmission job while elements 712 through 730 relate to a query command requesting status of a previously initiated transmission job.

[0069] Element 700 is operable in response to a command requesting the transmission of a new job object. Element 702 then writes data and parameters of the job in the persistent storage of the system to enable failure detection and recovery. User data files associated with the job object are copied to user data files 710. Other parameters required for failure detection and recovery are written to the transmit job database 708.

[0070] Element 704 is then operable to create the various objects 706 required for invocation of the send method. The send method is then invoked at element 711 to send the transmission job to the second secured computing system over the communication medium (via communication path 354 of figure 3).

[0071] Element 712 is operable to process a job object whose command requests a status update (the query method of the job). First, element 714 locates the identified job in the transmit job database 716. Element 718 then determines whether the database entry so located indicates that the job has been completed (fully acknowledged by the second secured computing system). If element 718 determines that the job has been fully acknowledged, element 720 is then operable to remove persistent data (708 and 710) related to the completed job. As discussed further herein below, in the preferred embodiment, completed job information is left in the job database for a period of time to permit "off-line" analysis and gathering of statistics. The period of time is preferably configurable by administrative users of the present invention. Processing then continues with element 726 to generate a query response indicating a completed status for the identified job.

[0072] If element 718 determines that the job is not yet completed, processing continues with element 726 to create an appropriate response to the query. The present status of the job is associated with a status object 728 created by element 726. Element 730 then returns the status object to the sending application which invoked the query method. The status object is returned via IPC communication path 352 of figure 3.

[0073] Figures 8, 9 and 10 provide additional data flow details regarding the operation of reception manager 306 of figure 3. A transport receiver process 800 of figure 8 is a process which is waiting for information to arrive from a transmission manager. In the preferred embodiment, process 800 is implemented as a small, stand alone application program launched in response to receipt of information. Those skilled in the art will recognize the equivalence of implementing such a process as a background task, a daemon, within the computing system. Such design choices are well known to those in the art.

[0074] As noted above, the communication medium may include network connections and associated network protocols such as email, FTP and HTTP networking protocols, direct modem connections, as well as manual transmission medium including magnetic or optical storage media. In whatever form the data is received from communication path 354, transport receiver process 800 writes received data into the INBOX directory structure 802 for subsequent processing by the reception manager. The receive office process 804 is an asynchronously operating process which reads INBOX directory structure 802 to detect when new transmission jobs have arrived. Header information associated with a received job identifies the particular transaction ID 808 associated with the job as transmitted to process 800. The transaction ID is then used in encoding the file name for files stored in the INBOX directory structure 802. The transaction

ID is also used to identify a particular unique OUTBOX directory structure 806 associated with the particular transaction. Information read from INBOX 802 by receive office process 804 is therefore written to an appropriate OUTBOX directory structure associated with the particular transaction ID 808 of the transmission. In this way, receive office process 804 is capable of processing a number of simultaneous transmission connections from a plurality of transmission managers.

5 An appropriate receiving application 810 is launched by processing of receive office 804 for purposes of receiving and further processing of a particular transmission job. Receiving application 810 reads the received information from its associated OUTBOX directory structure 806.

[0075] Figures 9 and 10 provide additional detail of the data flow within receive office process 804 of figure 8. Element 900 represents processing within the receive office process 804 which determines when a complete transmission for a particular transaction ID has been received in the INBOX directory structure 802. When element 900 recognizes that a complete transmission is presently wholly contained within INBOX directory structure 802, element 904 processes information in the INBOX directory structure 802 to gather and reconstruct the pieces of the job transmission in the proper order. The complete transmission 906 reconstructed in the appropriate sequence by processing of element 904 is then further processed by element 908 to decompress or decode any encoded information in the job and job objects 910 with associated data files are thereby created.

[0076] The "un-processed" transmitted job represented by job objects 910 is forwarded to element 912 for further processing to restore original filenames, if required, to the files transmitted in the job. The un-processed files with original filenames restored are then forwarded to element 914 which writes the files into the OUTBOX directory structure 806 associated with this particular transaction. The OUTBOX directory structure 806 serves as the persistent storage for data files associated with the job objects 910. Other persistent data associated with this particular job is then written by processing of element 918 into a job database file 920.

[0077] Element 922 is then operable to determine whether the sending application associated with this transaction ID is "registered" to initiate another receiving application process. The reception manager portion of the present invention maintains a database file recording sending processes which have initiated a transaction with a corresponding receiving process. An administrative configuration process within the second secured computing system generates a mapping table file which stores information to identify the required receiving process associated with a particular transaction.

[0078] If a receiving process is so registered to be launched, the process to be initiated is identified and initiated by processing of element 924. If no such application (receiving process) is to be initiated, no further processing is required as indicated by element 926.

[0079] Figure 10 is a data flow diagram providing additional details regarding the operation of receive office process 804 of figure 8 when responding to a request for acknowledgment from the receiving application/process (the ACK method of the reception manager). Element 1000 is first operable to receive the acknowledgment request (ACK method invocation) via IPC communication path 354 from the receiving process. An acknowledgment request includes as a parameter thereof a transaction ID identifying the particular job for which the acknowledgment request is solicited.

[0080] Element 1002 is next operable to check the job database file 920 to locate the transmission manager return address associated with the supplied transaction ID. The job database file 920 associated with the reception manager records transaction ID and associated control data required to complete a particular transmission transaction.

[0081] Element 1006 is then operable to prepare an appropriate acknowledgment response. Specifically an acknowledgment response is preferably formulated and transmitted in a manner which mirrors the transmission manager discussed hereinabove. Still more specifically, element 1006 prepares information for transmission to a reception manager resident in the other secured computing system. Element 1008 is then operable to invoke the send API method (as discussed above) to transmit the prepared acknowledgment back to the requesting transmission manager. The transmission of the acknowledgment message is performed in a manner identical to that of a job transaction in that the send method of the transmission manager portion of the invention is invoked. However, in transmission of an acknowledgment message, no acknowledgment therefore is required from the other side. This avoids an essentially infinite loop of acknowledgments "ping-ponging" between the two systems.

[0082] Figure 11 is a data flow diagram describing additional details of the operation of the receiving application 308 of figure 3. The receiving application preferably receives a transaction in the form of a command line to start the application (executed by a shell or command interpreter of the computing system) and parameters for operating the application. The parameters may, for example, identify the files of data transmitted from the sending application by their respective original file names in the OUTBOX directory structure path.

[0083] The receiving application then performs whatever application specific processing is required as depicted in element 1104. Element 1106 then creates a courier object to communicate with the reception manager. The ACK method of the courier object is then invoked by element 1108 to cause an acknowledgment message to be sent back to the sending process. The ACK method communicates the ACK request to the reception manager via IPC communication path 360.

[0084] Those skilled in the art will recognize that the receiving process may acknowledge the transmission from the

sending process at any time in its processing. In other words, the semantic interpretation of the ACK message is up to the design of the receiving process. The ACK message may be sent as soon as the receiving process is invoked to indicate merely that the job has been received. The ACK message may be sent following processing of the data by the receiving process to indicate that the requisite operations have been completed. Alternatively, the sending process
5 may indicate in job parameters that no acknowledgment is required for the transaction. In such a case, the receiving process need not generate an ACK message at all. Still another option for the communicating processes is to utilize the partial ACK method of the invention. The partial ACK method generates a different acknowledgment message which may be useful for indicating partial success of the transaction to the sending process. In the preferred embodiment, the receiving application determines when to send the ACK message and determines whether it is a normal ACK
10 or a partial ACK. The sending application simply awaits receipt of such ACK messages and interprets them accordingly as indicative of completion or only partial completion of the transaction, respectively.

METHODS OF THE PREFERRED EMBODIMENTS

15 [0085] Figure 14 is a flow chart describing methods of the present invention operable within a sending application. As noted herein above, a sending application invokes services of the present invention by calling functions defined in an application programming interface (API) of the present invention. The sending application specifically invokes API functions relating to the transmission manager portion of the present invention. As described further herein below, the interface between sending application, transmission manager, and other components of the present invention, are described in terms of object oriented design paradigms. Those skilled in the art will readily recognize other programming design paradigms (i.e., purely procedural paradigms, etc.) which may be utilized to implement the features of the present invention. Further, those skilled in the art will recognize the equivalence of implementations using custom circuit designs for various features of the present invention as compared to software designs operable on general-purpose computing devices.

20 [0086] At element 1400, the sending application first invokes an API function of the transmission manager to create a transmission manager courier object. The courier object is essentially the means by which the sending and receiving processes request services of the manager of the present invention. The API functions (described further herein below) are methods of this courier object. Element 1402 is next operable to invoke the send method of the courier object. A transaction ID is returned from the send method invocation for use in subsequent tracking of the progress of the job.
25 The transaction ID so returned is saved by operation of element 1404. Subsequent query operations to determine the status of the job use the saved transaction ID to identify the job for which status is requested. Next, elements 1406 and 1408 are iteratively operable to determine when the job is successfully transmitted to the destination receiving application. In particular, element 1406 is operable to invoke the query method of the courier object to determine the present status of the job as identified by the saved transaction ID associated with the job. Element 1408 is then operable
30 to determine whether the query method invocation results in a status indicative of job completion. If the job is completed, the processing of figure 14 is completed. If the transmission job is not yet completed, processing continues by looping back to element 1406 to repeat the query process.

35 [0087] Those skilled in the art will readily recognize that a number of jobs may be associated with a particular courier object. In other words a plurality of jobs may be directed to a single receiving application via a single transmission manager courier object (created by element 1400) and then iteratively invoking elements 1402 and 1404 to define a new job and transmit that job. Furthermore, those skilled in the art will recognize that the transmission, once initiated by processing of element 1404, continues asynchronously with respect to the sending application. The sending application may proceed with any other processing desired and invoke the query method (element 1406 and 1408) at any later time to poll for the present status of any pending transmission jobs.

40 [0088] Figure 15 is a flow chart describing methods of the present invention operable within the transmission manager for coordinating and managing robust, reliable exchange of data with a corresponding reception manager operable within a second secured computing system. The first flow chart (elements 1500 through 1508) describes the send method associated with courier objects of the transmission manager. The second flow chart (elements 1520 through 1530) describes the query method operable within the transmission manager courier objects.

45 [0089] Element 1500 is first operable to receive a send job request from the sending application process as described above with respect to figure 14. In response to receipt of such a job request, element 1502 is next operable to create a job object to control processing above this particular transaction job. Element 1502 is further operable to associate with the job object all required control information and copies of all supplied data files from the sending application process. Element 1503 then generates a unique ID (GUID) used as the transaction ID (also referred to herein as a job
50 ID). The job ID is used by all cooperating processes to identify a particular transaction or job in the job database.

55 [0090] Element 1504 is then operable to identify the preferred transmission medium associated with the identified destination receiving application. The database file associated with the transmission manager includes information identifying one or more transmission media appropriate for communicating with an identified destination. Element 1504

is therefore operable to select a preferred such medium. The preferred medium may be selected in accordance with performance standards of the various media or other standards applicable to particular environments. For example, it may be preferred initially to select a high-speed network connection available for transmission of data files associated with a particular transmission job. However, if the network communication link is unavailable for any reason, a slower manual transmission medium (i.e. magnetic or optical storage medium) may be selected as an alternate communication medium.

[0091] Element 1506 is then operable to invoke the transmission method for the identified destination. As noted herein above, the transmission method may include techniques appropriate to the particular medium for encoding and/or encrypting the job data files before transmission through the selected medium. For example, certain transmission media may require a large data file to be broken into multiple segments (i.e., some email protocols require large messages to be so recomposed into smaller segments). Or, for example, other transmission media may require that binary information be encoded into displayable ASCII text for transmission over the selected medium. All such encoding and/or encrypting is therefore preferably embodied within the processing of element 1506 as required for the particular selected medium.

[0092] Element 1508 is then operable to make the job object created by element 1502 persistent by storing the job object information in the transmission job database. Making the job object persistent enables failure detection and recovery techniques to be applied to the transmission of the job. Finally, element 1509 is operable to return the generated transaction ID to the calling process for later use in determining the status of the job.

[0093] Elements 1520 through 1530 describe the query method of the transmission manager of the present invention. Specifically, element 1520 is first operable to look up the current status of a pending job as presently recorded in the transmission job database. Element 1522 is next operable to determine whether a record identifying the job is so located within the job database. If the job is no longer recorded in the job database (i.e., the job has been previously completed and removed from the database), element 1530 is next operable to return an error job status to the sending process thereby completing processing for the method. The "job no in database" error status so returned informs the calling process that the job has been completed and thence removed from the database or that the transaction ID used does not correspond to any known job.

[0094] If element 1522 locates the identified job in the transmission job database, element 1526 is next operable to determine whether the present status as recorded in the job database indicates that the job has completed. If the job has been completed, element 1528 is next operable to remove the job entry from the transmission job database. In either case, processing then continues with element 1530 to return the present job status (i.e., complete or not yet complete) to the calling (sending) application process thereby completing processing of the method.

[0095] As noted herein above, in the preferred embodiment, completed jobs remain in the job database for a period of time following completion to permit "off-line" statistical gathering and analysis. A cleanup timeout period is associated with the completed job entry in the database such that time out processing methods (described below) will eventually delete the completed job entry from the job database.

[0096] Figure 16 is a flow chart describing daemon (background) processes operable within the transmission manager of the present invention. In particular, elements 1600 through 1618 describe the processing of the timeout monitor daemon operable to detect a timeout condition in a pending or completed transmission job. Elements 1620 through 1624 described a method, asynchronously initiated, to detect and process the receipt of an acknowledge message from the reception manager in the second secured computing system.

[0097] Element 1600 is first operable to initialize a local pointer variable (JOB) to point to the first pending job in the transmission job database. Elements 1602 through 1616 are then iteratively operable to determine if the job presently pointed to by the local pointer variable has timed out waiting for a timed event. Specifically, element 1602 first determines whether the job has a completed status. If not, element 1604 next determines whether a predetermined retransmit timeout period associated with the job has expired. As noted above, the retransmit timeout period for each job may vary depending upon the selected transmission medium and other parameters of the transmission manager. If element 1604 determines that the timeout period has expired, elements 1606 in 1608 are then operable to re-transmit the job. Specifically, element 1606 initiates retransmission of the job for which the retransmit timeout period has expired. Element 1608 then resets the retransmit timeout period for the job in view of the initiation of retransmission of the job. Processing then continues with element 1614. If element 1604 determined that the retransmit timer has not expired, processing continues with element 1614.

[0098] If element 1602 determines that the job status is complete, element 1610 is next operable to determine whether the cleanup timeout period has expired for the job. As noted above, in the preferred embodiment, completed jobs remain in the job database for a period of time (the cleanup period) to permit post processing statistical analysis or other analysis or reporting functions. The cleanup period is preferably configurable to permit a system administrator to determine the needs of the installation. If the cleanup period has expired, element 1612 is operable to remove the completed job record from the job database. Processing then continues with element 1614.

[0099] Processing then continues with element 1614 to determine whether additional job records need be processed

in the transmission job datacase. If more job records are present in the transmission job database, element 1616 is next operable to set the local pointer variable to a next job entry in the job database. Processing then continues by looping back to element 1602 to inspect the status for this next job. If element 1614 determines no further jobs are pending in the transmission job database, element 1618 is operable to delay a predetermined period of time (i.e., a tick of the timeout monitoring clock). Processing then continues by looping back to elements 1600 to review the timeout status of all remaining job records in the transmission job database.

[0100] Elements 1620 and 1622 are operable as a daemon process within transmission manager to recognize and process the asynchronous receipt of an acknowledgment message from the reception manager operable in the second secured computing system. Specifically, element 1620 is operable to receive the asynchronously transmitted acknowledgment message (ACK). Element 1622 is then operable to update the completion status for the corresponding transmission job in the transmission job database in accordance with the newly received acknowledgment.

[0101] Figure 17 depicts a flow chart of a method operable within the reception manager in the second secured computing system to receive transmitted data from the transmission manager operable in the first secured computing system.

[0102] Element 1700 is operable to receive a job transmission from the transmission manager. As noted above, depending upon the selected transmission medium, a single transmission job may be decomposed into a plurality of segments. Element 1700 is therefore operable to receive each such segment of a job. Element 1700 is further operable to store all such received segments in an INBOX structure for temporary retention until all segments have been received. The INBOX structure is a temporary heap storage or repository for holding the received segments of transmitted jobs.

20 Each segment is stored in a named file, the name of which includes the transaction ID of the job to associate it with other segments of a particular transmission job.

[0103] Element 1702 is then operable to determine whether all segments of a particular job transmission have been received. If element 1702 determines that no complete job has yet been received and stored in the INBOX segment repository, processing continues by looping back to element 1700 to await reception of a complete transmission job.

25 [0104] If element 1702 determines that at least one complete transmission job presently resides within the INBOX repository, element 1704 is next operable to "unprocess" the received transmission job. As used herein, "unprocessing" preferably includes decoding of any encoded information, decryption of any encrypted information, and the combining related segments into their original files as defined by the transmission job parameters. As noted above, transmission jobs may be associated not only with a plurality of segments, but the plurality of segments may constitute a plurality of files associated with the transmission job. Element 1704 is therefore generally responsible for re-constructing the originally configured files as associated with the transmission job by the sending application process.

30 [0105] Element 1708 is then operable to create a job object as a repository of information regarding the completed transmission job. Element 1710 is next operable to associate the receiving application of the received transmission job with the created job object. Element 1712 is then operable to make the created job object persistent by storing it in the job database associated with the reception manager. So storing the received job object in a persistent received job database permits failure detection and recovery techniques within reception manager as described herein.

35 [0106] Element 1714 is then operable to store the unprocessed job (reconstructed into its original file structures) in an OUTBOX directory structure corresponding to the transaction ID of the transmission job. Each transmission job received and unprocessed by elements 1700 through 1704 is preferably stored in a unique OUTBOX directory structure associated with that specific transaction.

40 [0107] Element 1716 is then operable to locate the receiving application process corresponding to this transmission job. As noted above, parameters of the job preferably include identification of a particular receiving process or application intended to further process the received job. Element 1718 is then operable to determine whether an application and has been so identified as associated with the transmission job. If none has been selected, processing continues by looping back to element 1700 to await reception of further segments associated with further transmission jobs.

45 [0108] If a receiving application or process has been identified and located in association with the received job transmission, element 1720 is then operable to launch the identified receiving application to process the received job data. The receiving application is launched with a command line and associated parameters identified as job parameters and associated with the transmission job. Processing then continues by looping back to element 1700 to await reception of further segments of other transmission jobs.

50 [0109] Figure 18 depicts flow charts of methods of the present invention operable within the reception manager in the second secured computing system. Element 1802 through 1806 describe an acknowledge method operable within reception manager in response to invocation by receiving application following completion of its processing of the received transmission job. Element 1802 is operable to send an ACK message originating transmission manager.

55 [0110] As noted above, acknowledgment aspects of the present invention are substantially symmetric with respect to the transmission manager and reception manager of the present invention. In this case, the acknowledgment request is generated by the receiving application to indicate to the transmission manager that the transmitted job has been received and processed. The transmission of the ACK message uses the same objects and methods discussed above

with respect to the transmission manager. As noted above, the transmission manager and reception manager are preferably a single module providing the API functions of both management features in both systems. The acknowledge method therefore sends its ACK message in essentially the same way the sending application sends the original transaction - by invoking the send method of a courier object. In the case of the acknowledge method transmission, the transmitted message is sent without requesting an acknowledgment from the other side. Parameters of the job do not request an acknowledgment. This prevents an "infinite loop" of acknowledgments causing further acknowledgments, etc.

5 [0111] Element 1806 is next operable to remove the job from the receive job database. Having completed all processing of the received transmission job, no persistent storage is required relating to the particular received job transmission.

10 As noted above, completed jobs are "removed" from the job database after a configurable timeout period to permit statistics gathering and other analysis of the communication paths.

15 [0112] Elements 1820 through 1823 are operable within the identified receiving application or process designated to process the particular received job transmission. Specifically, element 1820 is operable to perform whatever application-specific processing is required for the received job transmission data files. Element 1822 is then operable to invoke the acknowledge method of the reception manager as discussed above. The acknowledge message so generated verifies reception and processing of the job data to the originating (sending) process. As noted above, those skilled in the art will recognize that the receiving process may defer the generation of the acknowledge to any point in its processing deemed appropriate for the transaction. The acknowledge may be generated as soon as the data is received or after any point in the application specific processing of the received data.

20 [0113] Lastly, element 1823 is operable within the receiving application to remove the job data files from the OUTBOX directory structure created by the reception manager. The receiving application may delete this directory structure at its discretion or, if appropriate, leave the structure for subsequent processing following the ACK message returned to the sending application. Well known "garbage collection" techniques may be employed by the reception manager of the present invention to subsequently remove OUTBOX directory structures inadvertently left in the system by ill-behaved receiving application processes.

25

API DEFINITIONS OF THE PREFERRED EMBODIMENTS

[0114] As noted above, the transmission manager and reception manager of the present invention are preferably implemented as object oriented API modules to be invoked by sending and receiving applications desiring to communicate between secured systems. The following API definitions represent the best presently known mode of practicing the invention.

30 [0115] The features of the present invention described above with respect to the transmission manager and the reception manager are provided as an API linked with the sending and receiving applications. The management features are preferably implemented in conjunction with lower level protocol control modules referred to as Transport Office. The management features are therefore referred to as Transport Office Manager (or by the acronym TOM). The API provides an object oriented interface to the features of the present invention. The sending and receiving applications are therefore linked with the API library. Desired objects are instantiated by the applications and methods associated with those objects are invoked to perform the desired transaction management.

35

TomIL Classes: An Overview

[0116] The classes described in this document represent the complete external interface for the TOM Interface Library (TomIL). The TomIL is itself an interface for communicating with the Transport Office Manager (TOM). Users of the TomIL instantiate TomIL objects representing transport protocols, transport policies, and jobs, and then pass those objects to the communications hub of TomIL: CTomilCourier. The CTomilCourier object then communicates directly with TOM via sockets. Once transactions have been sent via the CTomilCourier, those transactions can be queried for status, thus creating other TomIL objects.

40 [0117] TomIL objects all share the basic characteristic that they are handled as pointers at the client level. There is no way to construct any TomIL object by simply using the object's constructor. Since objects are all accessed at the client level through pointers, that means no copy constructors or assignment operators have been provided for TomIL objects exposed to clients. There has been provided, however, a virtual copy constructor for all objects that are not treated as Singletons(127). The notation Singleton(127) refers to the singleton design pattern described in the book: Design Patterns by Gamma, Helm, Johnson, and Vlissides. By convention, Prototype(117) would refer to the prototype design pattern in the same text.

45 [0118] There are three basic types of TomIL objects:

Type I Objects: These objects are objects that are created via the static "Create()" method accompanying the object's class or via a return value from a CTomilCourier method, such that when the object is created, a full copy of

the object comes into existence and ownership of that object passes from TomIL to the client. This means that the client is responsible for deleting the memory associated with said object once the object is no longer useful
For example:

```
5      CTomilPolicy* policy = CTomilPolicy::Create(10, 10, 5, TOMIL_ACK );
// Use the object here to do whatever
delete policy;
```

10 OR

[0119]

```
15      CTomilTransactionID* TransactionID;
CTomilCourier* Courier = CTomilCourier::Instance();
TransactionID = Courier->Send( job, app );
// Use the transactionID object here to do whatever
delete TransactionID;
```

[0120] Type I Objects all have virtual copy constructors associated with them. This is how clients can make full, deep copies of an object given that the client simply has a pointer to the object. All virtual copy constructors are called "clone ()" and always take no parameters.

For example:

```
30      CTomilPolicy* policy1 = CTomilPolicy::Create( 10, 10, 5, TOMIL_ACK );
CTomilPolicy* policy2;
// Create a full, deep copy of the object policy1 which is now pointed to by the
// pointer "policy2".
// These objects are fully independent of each other and each can be deleted
35      // without affecting the other.
policy2 = policy1->clone();
// Make sure to also delete policy2 once you're finished with it or you will create
// a memory leak.
```

40 [0121] Type I Objects include:

- CTomilPolicy: create via CTomilPolicy::Create() or clone()
- CTomilJob: create via CTomilJob::Create() or clone()
- CTomilApplication: create via CTomilApplication::Create() or clone()
- CTomilAddress: create via CTomilAddressFactory's or CTomilAddressStrategy's or clone()
- CTomilTransactionIDIterator: create via CTomilcourier::GetAllIDs() or clone()
- CTomilTransactionID: create via CTomilTransactionID::Create() or clone()
- CTomilAddressStrategy_SpecificProtocol: create via CTomilAddressStrategy_SpecificProtocol::Create() or clone ()
- CTomilTransactionStatus: create via CTomilCourier::Query() or clone()
- CTomilVersion: create via CTomilVersion::Create() or clone()
- CTomilError: created via "catch" statements or clone()

55 Type II Objects: These objects are referred to as Singleton(127) objects. A Singleton(127) object ensures that a class has only one instance and provides a global point of access to it. Objects that are Singletons(127) are represented as pointers throughout a user's code, and the pointer to a Singleton(127) is delivered through its Instance() method. Pointers to Singletons(127) are owned by the Singleton(127) and will be disposed of properly when the application

using the Singleton(127) exists. No clean-up intervention is needed on the part of the user of the Singleton(127). As a matter of fact, clients must explicitly NOT try to delete any pointers which point to Singletons(127)

For example:

```

5   {
10    CTomilCourier* courier = CTomilCourier::Instance();
15    courier->Send( job, app );
20    } // local scope

// At this point in the code the CTomilCourier object still exists and may be
// referenced again by simply calling its ::Instance() method. The object will
// be destroyed automatically when the entire application exits.
25

```

[0122] TomIL Type II objects include:

- CTomilAddressFactory (abstract base class)
- CTomilAddressFactoryEmail
- CTomilCourier

Type III Objects: These objects look and act pretty much like Type I objects: they have a clone() method and are referred to via pointers, but they are instantiated by factory or strategy methods because users of these objects don't have enough information exposed to them to construct the objects themselves.

For example:

```

30   CTomilAddress* address;
35   CTomilAddressStrategy_SpecificProtocol* strategy =
        CTomilAddressStrategy_SpecificProtocol::Create();
        address = strategy->CreateAddress( "atl" );
        // Use the address object here for whatever...
        delete address;

```

[0123] TomIL Type III objects include:

- CTomilAddress is the only Type III object in the TomIL.

TomIL Error Strategy

[0124] The error strategy throughout TomIL is for methods to throw CTomilError objects whenever exceptional conditions are encountered, however. TomIL constructors will never throw errors. Any TomIL method exposed to the user that can throw a CTomilError object will be documented as such. It is strongly recommended that any TomIL method that can throw an error be surrounded with Try/Catch code whenever it is called. Always catch CTomilError objects by pointer.

For example:

50

55

```

try
{
    strategy->CreateAddress( "bad_address" );
}
catch ( CTomILError* e )
{
    // process error here
    delete e;
}

```

TomIL Address Books

[0125] An address book is a file containing a textual collection of addresses and policies that can be used by TomIL. Addresses and policies are spelled out in their entirety and then given a short-hand alias by which they can be referred. Address books are completely optional when using TomIL. They merely provide a short-hand access to well-known and often used addresses and policies. TomIL clients are welcome to construct their own addresses and policies at run time using TomIL objects.

[0126] For example, a short-hand alias for the Atlanta Response Center might be the string "ATL", and its full email address might be respCenter@atlhub.atl.hp.com. A short-hand alias for the default policy used to send email to the address ATL might be the string "DefaultPolicy" and the policy parameters might be (60;5;60;1).

[0127] Address books are preferably human-readable text and can be created with any textual editor. A standard text editor or simple web-based address book construction tool may be used to construct the file.

TomIL address book syntax

[0128] A TomIL address book can be named any legal file name for the operating system you're using, but all TomIL address books should have the suffix ".abk" so they can be easily identified as TomIL address books.

[0129] Address book comments are preferably full-line comments only and may be placed on any full line of the address book. Comments are denoted by having a "#" character as the first character on the line.
For example:
This is a comment line.

Comments are preferably not embedded within functional lines in your address book. Blank lines: Blank lines may be inserted anywhere in the address book. All blank lines will be ignored. Delimiters: The field delimiter in a TomIL address book is always the character ";" (semi-colon). It is not necessary to end lines in an address book with a ";". The end-of-line character is plenty to denote the end of a field. Spaces: Do NOT embed spaces in address headers, address names, or policy names. Section headers: The format for a section header is "**H;name**".

For example:
This is the header for the Atlanta section
H:ATL
Addresses: An address section is delimited by the line "--Addresses--"
For example:
--Addresses--
Individual addresses (Email): Each email address has the format: "Email;address@address"
For example:

Sample email address. Any legal email address can follow the ";"
on the following line
Email;rc@atl.hp.com

Policies A policy section is delimited by the line "--Policies--".
For example:
--Policies-
Individual policies: Each policy has the format:

"PolicyName;TimeOut; NumRetries TBR;Ack"
 For example:

5 # The following policy is named "DefaultEmail" and has a TimeOut
 # value of 60, a Number-of-retries value of 5, a time-between-retries
 # value of 30, and is acknowledged (1 instead of 0)
 DefaultEmail;60;5;30;1

10 [0130] An address book preferably has exactly one Email address defined for each section header, and as many uniquely named policies as desired for the same header. An address book may have as many uniquely named sections as needed.

15 Sample TomIL address book

[0131] The following is a valid sample TomIL address book.

20 #
 # This is a Sample TomIL Address Book
 #
 #####
 # First-choice header for Atlanta
 #####
 H:Atl-primary
 --Addresses--
 Email;rc@atl.hp.com
 --Policies--

35 LargeEmail;500;4;300;1
 MediumEmail;50;4;30;1
 SmallEmail;5;4;3;1
 #####
 # This is a header for an alternate Atlanta address
 #####
 H:Atl-secondary
 --Addresses--
 Email;rc_secondary@atl.hp.com
 --Policies--
 LargeEmail;300;2;30;1
 #
 # etc....
 #

TomIL Use Model

55 [0132] The TomIL is a simple library to use. Basic C++/OO (object oriented) paradigms are followed in using the TomIL with an application process (sending or receiving application). Certain objects in TomIL are created on the heap. It will be noted below which of these heap objects belong to TomIL and which objects have their ownership transferred back to the client. The calling process is responsible for deleting all heap objects that TomIL transfers to the client

process.

Overview

5 [0133] The basic function of the TomIL is to gather data and parameters, packaging them into objects that funnel through the communication gateway to TOM, namely the CTomilCourier object. The following sequence is the same basic sequence all client applications will follow to communicate with TOM via TomIL. Each communication with TOM is referred to as a "transaction", and each transaction can contain one or more files to be transmitted.

10 1) The application *registers* itself with TomIL by creating a CTomilApplication object. This object is then passed to CTomilCourier methods (Send(), GetAllIDs()).
 2) The application needs to decide if a TomIL address book will be used. There is no problem inter-mixing the use of run-time address/policies with those retrieved from address books, but if an address book is to be used, then its location needs to be set as soon as possible in the application if the application is not going to use the default address book location. Static methods on the CTomilApplication object are used for this purpose. Address book locations can be set and reset anywhere in the application, but the client must make sure to have the proper address book location set before using any CTomil object that takes an address book alias as a parameter. If the address book location is in error, then CTomilError objects will be thrown when any method that uses an address book alias is called.
 15 3) The application must create a CTomilPolicy object to use for its transaction. This object can be created at run-time by passing exact parameters to the CTomilPolicy::Create() method, or by fetching parameters from an address book by using the CreatePolicyFromAB("alias", "alias") method. Either way is equally valid. Neither is preferred over the other.
 20 4) The application must create a CTomilAddress object to use for its transaction. Creating this object is the single most potentially confusing part of using the TomIL. CTomilAddress objects are not distinguished by protocol at the client level (i.e., an email address object looks just like an FTP address object at the client level). It's the internal implementations of these objects that distinguish them. For the first release of TomIL, there is only one supported protocol: email, but there will be others in future releases. A great deal of effort has been designed into TomIL to shield clients from differences incurred by different transport protocols. Email addresses look wildly different from FTP addresses, and both are completely different from HTTP addresses. For this reason, clients cannot directly create generic CTomilAddress objects with protocol-specific implementations. Factories or strategies are needed to accomplish this.

25 35 A CTomilAddressFactory is a low-level way to create a CTomilAddress object. In particular, a CTomilAddressFactoryEmail can be instantiated and then its CreateAddressFromAB("alias") method or CreateEmailAddress("actual_address") method can be invoked to create a CTomilAddress with an email-specific implementation. The latter of these two methods is the only way to create a CTomilAddress object when the application has an actual email address. Using a CTomilAddressFactory to create a CTomilAddress object is not recommended because it ties client applications to specific TomIL address protocols when the client is compiled. Client applications will be much better served to relinquish all address protocol decision matters to TomIL.

40 45 Using a CTomilAddressStrategy is the preferred way to construct CTomilAddress objects. Strategy objects in general encapsulate algorithms allowing them to vary independently from the clients that use them. An address strategy is an algorithm that computes the best possible address type for the client to use given the constraints passed to the strategy. For the first release of TomIL, there is only one address strategy, and it is extremely simple: use email protocol. The CreateAddress("alias") method on the CTomilAddressStrategy_SpecificProtocol object is how a CTomilAddress object is created. If client applications get in the habit of using address strategies instead of specific address factories, then they will remain completely divorced from any specific transport protocol that TomIL can add in future releases, thus allowing the client application to always be poised to take advantage of new protocol enhancements made to strategies without ever having to re-compile.

50 55 5) After CTomilAddress and CTomilPolicy objects have been created, they are passed to a CTomilJob object. It is through the AddFile() method of this object that files are grouped into a single job/transaction, all being sent according to the policies of the specific CTomilPolicy object to the address of the specific CTomilAddress object.

6) Once a CTomilJob has been created, that job is submitted to TOM via the Send(job, app) method on a CTomilCourier object. The Send() method will return a new CTomilTransactionID object. This object identifies the transaction requested of TOM and is the identifying object used for future queries of TOM regarding this specific transaction. This object may be passed to the Query(ID) method of a CTomilCourier which will return a CTomilTransactionStatus object that relays the status of the transaction specified by the given CTomilTransactionID object. The CTomilTransactionStatus object has many different Get() methods which allow clients to query different aspects of a given transaction.

Sample Code

[0134] There are 2 sample pieces of code that follow. The first is a valid sample TomIL client application. It follows the steps listed above in the section "Overview". The second is a sample Tom receiver.

5

10

15

20

25

30

35

40

45

50

55

```

#include <string.h>
#include <iostream.h>
#include <CTomilApplication.h>
#include <CTomilPolicy.h>
#include <CTomilAddress.h>
#include <CTomilAddressFactoryEmail.h>
#include <CTomilAddressStrategy_SpecificProtocol.h>
#include <CTomilJob.h>
#include <CTomilTransactionID.h>
#include <CTomilCourier.h>
#include <CTomilTransactionStatus.h>
#include <CTomilTransactionIDIterator.h>
#include <CTomilError.h>

// NOTE: The name of the address book we're using is passed into this
//       program. The path is specific to which ever platform we're
//       on: either NT or UX.
void main( int argc, char* argv[] )
{
    char* ab; // The path of the address book.
    ab = new char[ strlen( argv[1] ) + 1 ];
    strcpy( ab, argv[1] );

    //
    // Step 1 as written in the ERS: register this application.
    //
    CTomilApplication* app = CTomilApplication::Create( 'myApp',
                                                       "myVersion",
                                                       "myID" );

    //
    // (optional) Step 2 as written in the ERS: set the location
    // of an address book.
    try
    {
        CTomilApplication::SetAddressBookName( ab );
    } // try
    catch ( CTomilError* e )
    {
        cout << e << endl;
        delete e;
    } // catch

    //
    // Step 3 as written in the ERS: create a CTomilPolicy
    //           method 1, direct parameters
    CTomilPolicy* policy1 = CTomilPolicy::Create( 60, 5, 30, TOMIL_ACK );
    //
}

```

```

// Step 3 as written in the ERS: create a CTomilPolicy
//           method 2, use CreatePolicyFromAB()
CTomilPolicy* policy2 = 0;
5          try
{
    policy2 = CTomilPolicy::CreatePolicyFromAB( "Atl-primary",
                                              "LargeEmail" );
}
10         } // try
catch ( CTomilError* e )
{
    cout << e << endl;
15         delete e;
} // catch

// Since policy2 isn't used except for demo purposes, delete it.
20         delete policy2;
//
// Step 4 as written in the ERS: create a CTomilAddress object
//           method 1: use an actual address and
//           factory
//           NOT RECOMMENDED.
25         //
CTomilAddress* address1;
CTomilAddressFactoryEmail* emailFactory =
30             CTomilAddressFactoryEmail::Instance();
address1 = emailFactory->CreateEmailAddress( "myaddr@atl.hp.com" );

// Since address1 isn't used except for demo purposes, delete it.
35         delete address1;

//
// Step 4 as written in the ERS: create a CTomilAddress object
//           method 2a: use a CTomilStrategy.
//           This method is still not recommended
//           as it passes a specific address
//           factory to the strategy.
40         //
45         // Notice the subtle difference here in the CTomilAddressFactory
// pointer. The previous example had a specific
// CTomilAddressFactoryEmail pointer, this example MUST use a generic
// CTomilAddressFactory pointer because it's being passed to a
// CTomilStrategy.
50         CTomilAddressFactory* emailFactory2 =
            CTomilAddressFactoryEmail::Instance();
CTomilAddressStrategy_SpecificProtocol* addressStrategy;
55         addressStrategy =
            CTomilAddressStrategy_SpecificProtocol::Create( emailFactory2 );

```

```

CTomilAddress* address2,
try
{
    address2 = addressStrategy->CreateAddress( "AtI-secondary" );
} // try
catch ( CTomilError* e )
{
    cout << e << endl;
    delete e;
} // catch

// Since addressStrategy and address2 aren't used except for demo
15 // purposes, delete them.
delete addressStrategy;
delete address2;
//
20 // Step 4 as written in the ERS: create a CTomilAddress object
// method 2b: use a CTomilStrategy.
// This is the RECOMMENDED method to do
// this. Use the default constructor
// for the CtomilStrategy and you free
25 // the client from specific protocols.
//
CTomilAddressStrategy_SpecificProtocol* addressStrategy2;
addressStrategy2 = CTomilAddressStrategy_SpecificProtocol::Create();
CTomilAddress* address3;
30 try
{
    address3 = addressStrategy2->CreateAddress( "AtI-primary" );
} // try
catch ( CTomilError* e )
{
    cout << e << endl;
    delete e;
} // catch

// Since we're finished with addressStrategy2, delete it.
35 delete addressStrategy2;

//
40 // Step 5 as written in the ERS: create a CTomilJob object
// then add files to the Job
//
50 // These file names are hard-coded simply as examples of files that
// exist. Client code should not look like this other than to take
// into consideration the differences between UX and NT file names.
#ifndef _WIN32
    char* file = "c:\\winnt\\system32\\ansi.sys";
#else

```

```

char* file = "/usr/sbin/cron";
#endif

5      CTomilJob* job = CTomilJob::Create( address3, policy1 );

// Add a file to the job  Call AddFile() as many times as you need.
try
10     {
job->AddFile( file );
    } // try
catch ( CTomilError* e )
{
15     cout << e << endl;
    delete e;
} // catch

// We're finished with address3 and policy1, delete them.
20      delete address3;
delete policy1;

// Step 6 as written in the ERS: create a CTomilCourier object
// then use the Send(job) method.
//
25      CTomilTransactionID* TransactionID;

// Instantiate a CTomilCourier.
30      CTomilCourier* Courier = CTomilCourier::Instance();
try
{
35      TransactionID = Courier->Send( job, app );
    } //try
catch( CTomilError* e )
{
40      cout << e << endl;
    delete e;
} // catch

// We're finished with our job at this point, so delete it.
45      delete job;

// Step 6 as written in the ERS: query a transaction
//
50      CTomilTransactionStatus* status;
status = Courier->Query( TransactionID );

time_t timeSent = status->GetTimeSent();
55      cout << "Time sent = " << ctime( &timeSent ) << endl;

```

```

// Now get the policy parameters for this particular TransactionID.
CTomilPolicy* policy;
policy = status->CreatePolicy();
5   if ( policy != NULL )
    cout << "Number of retries = " << policy->GetNumRetries() << endl;
else
    throw "some kind of error";
10

if ( status->GetCurrentTry() !=
      CTomilTransactionStatus::GetUnknown() )
    cout << "Current Try = " << status->GetCurrentTry() << endl;
15

// Make sure to clean up objects when they are no longer useful...
delete policy;
delete TransactionID;
20   delete status;

//
// Bonus code: query the status of all your transactions
//           use a CTomilTransactionIDIterator
25
//
CTomilTransactionIDIterator* IDIt;
CTomilTransactionStatus* status2 = 0;

30   IDIt = Courier->GetAllIDs( app );
IDIt->Reset();

CTomilTransactionID* tid;
35   while( tid = IDIt->Next() )
{
    delete status2; // Prevent memory leaks.
    status2 = Courier->Query( tid );
    cout << "Status for job: (" << tid->GetID() << ") = "
        << status2 << endl;
40   } // for

45   delete IDIt;
delete status2;
// Don't delete tid because it just received a pointer assignment;
// there was never a new object created.
50   delete app;
} // main

```

55 [0135] This is sample receiver code:

```

#include <iostream.h>
#include <stdio.h>
#if !defined( _WIN32 )
    #include <unistd.h>
#endif
#include <CTomilCourier.h>
#include <CTomilTransactionID.h>
#include <CTomilError.h>

// Arguments to this program are : guid. file [.file[.file...]]
void main( int argc, char* argv[] )
{
    char* guidIn;
    if ( argc < 2 )
    {
        cout << "Error: bad arguments!" << endl;
        exit( 1 );
    } // if
    guidIn = argv[1];
    cout << "Guid value passed in = " << guidIn << "" << endl;
    // Instantiate a CTomilCourier.
    CTomilCourier* courier = CTomilCourier::Instance();
    // Create a valid transactionID object from the guid passed in.
    CTomilTransactionID* tid = CTomilTransactionID::Create( guidIn );
    // Loop through all the files passed in and validate each one:
    // check for existence.
    int arg = 3; // This is where the file arguments start
    int valid_files = 0;
    FILE* fp;

    while ( arg <= argc )
    {
        cout << "File: " << argv[ arg - 1 ] << " ";
        if ( ( fp = fopen( argv[ arg - 1 ], "r" ) ) == NULL )
            cout << "does not exist!" << endl;
        else
        {
            fclose( fp );
            valid_files++;
            cout << "exists." << endl;
        } // else

        arg++;
    } // while
}

```

```

char message[1024];

5   sprintf( message, "Data files sent: %d. Data files received: %d",
            argc - 2, valid_files );
cout << message << endl;
// Now that we've validated the files, send a partial ack back for
// the guid.
10 // Partial acknowledgments are optional.
try
{
    courier->PartialAck( tid, "Partial Ack. Processes pending." );
}
15 catch ( CTomilError* e )
{
    cout << e << endl;
    delete e;
20    } // catch
// Do whatever processing your app needs to do here.
// Now send a full acknowledgment.
try
{
    courier->Ack( tid, message );
}
25 catch ( CTomilError* e )
{
    cout << e << endl;
    delete e;
}
30    } // catch
35 delete tid;
} // main

```

TomIL Application

40 [0136] The TomIL Application is an object that allows an application to "register" itself with TomIL. The application registers its name, version, and unique ID (if necessary) then passes this information along to TOM via the CTomil-Courier methods Send() and GetAllIDs(). This application information is used on the receiving side to identify what programs to start and where to put transmitted data or which transactionIDs to return.

Class --> CTomilApplication

45 [0137] The constructor for CTomilApplication is protected to force instantiation through the static CTomilApplication::Create() method.

Public Methods:

50 [0138] static CTomilApplication* Create(const char* appName,
 const char* appVer,
 const char* appID = 0)

[0139] This is how clients create a CTomilApplication object. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

5 appName : The name of the client application. Must not be null, o/w a CTomilError will be thrown.
 appVer: The version of the client application Must not be null, o/w a CTomilError will be thrown
 appID : The id of this particular invocation of the client application. This parameter is only necessary if multiple invocations of the application will be running at the same time. Important!!! - This id must be persistent and unique.
 10 That is to say that if this particular client invocation dies unexpectedly and then restarts, it must come up with the same id as it had before it crashed in order for this parameter to be useful. This parameter is completely optional and will probably be left blank in most all applications.

10 NOTE: Regarding the appName, appVer, and appID parameters, all embedded spaces will silently be replaced with the "_" (underscore) character. If any of the parameters contains characters other than (a-z, A-Z, 0-9, .. _) a CTomilError will be thrown.

CTomilApplication* clone() const

15 [0140] This method is how a client creates a deep copy of this object. For example: CTomilApplication* app2 = app1->clone();

const char* GetName() const

20 [0141] This method returns the name that was used to create this CTomilApplication object. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

const char* GetVer() const

25 [0142] This method returns the version that was used to create this CTomilApplication object. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

const char* GetID() const

30 [0143] This method returns the ID that was used to create this CTomilApplication object. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

static void SetAddressBookName(const char* newName)

35 [0144] Clients use this method to set the location of the address book they want to use. Once the address book location is set, it will be in effect for the entire client application until another call is made to SetAddressBookName() or a call is made to RestoreDefaultAddressBook(). SetAddressBookName() need never be called. If this method is never called, the default TOMIL address book will be used. See RestoreDefaultAddressBook() for its location. This method will throw a CTomilError object if the address book specified by "newName" is invalid.

40 newName: This is the fully qualified path name to an alternate address book the client wants to use. This path name must be of the appropriate type for the operating system on which the client's code is running. For example: /var/opt/foo/book.abk on HP-UX, or C:\myPath\foo\book.abk on NT. The suffix ".abk" is not necessary, but helps to identify (to the OS or an observer) the file as an address book.

45 static const char* GetAddressBookName()

[0145] Clients use this method to get the name of the current address book being used throughout the client application. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

50 static void RestoreDefaultAddressBook()

[0146] Clients use this method to restore the current address book location being used throughout the client application to its original location. The default location of the TomIL address book is: /var/opt/TOM/lib/Tom.abk on HP-UX or C:\Tom\lib\Tom.abk on NT. This method will throw a CTomilError object if the default address book doesn't exist or is invalid.

TomIL Version

[0147] The TomIL version object allows clients to determine the current version of the TomIL that they are using.

5 Class --> CTomilVersion

[0148] The constructor for CTomlVersion is protected to force instantiation through the static CTomlVersion::Create() method.

10 Public Methods:

```
static CTomilVersion* Create()
```

[0149] This is how clients create a CTornilVersion object. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

CTomlVersion* clone() const

20 [0150] This method is how a client creates a deep copy of this object. For example: CTomlVersion* ver2 = ver1->clone();

```
const char* Get ( void )
```

[0151] This method returns a character string representing the current version of TomIL being used. For the first release of TomIL, this string will always be "A.01.00". The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

TomIL Policy

[0152] A policy is a packaging of all parameters that define attributes of how jobs can be sent using a CTomilCourier. A CTomilPolicy contains all of the transmission management variables for a particular job. Things like how long to wait before timing out, how many times to retry the transmission before generating an error, how long to wait before attempting a retry, and whether or not an acknowledgment is requested are contained in the policy class.

35

[0153] Policy variables can either be set explicitly in the CTomilPolicy::Create() method, or by passing a token to a method that enables a look-up of the corresponding values from a client's Address Book. The constructor for CTomil-

Policy is protected to force instantiation through the static `TT_WAKE_UP(TOMIL_NACK_TOMIL_ACK)`

Part II: Mathematics

[0155] This is how clients create a CTomlPolicy object. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

TimeOut: This tells TOM how long to wait for an acknowledgment before it considers a transmission to be timed out. It has a unit of "minutes" and will be limited to (0 <= TimeOut <= 14400) (10 days)

NumRetries: This tells TOM how many re-transmissions to make on a job should any of the previous transmission fail or time out. It is unit-less and will be limited to the range ($0 \leq \text{NumRetries} \leq 20$)

TBR : This tells TOM how long to wait between re-transmissions (Time Between Retries). It has a unit of "minutes" and will be limited to the range (0 <= TBR <= 2880) (2 days)

Ack: This tells TOM whether or not to expect an acknowledgment for a particular transmission.

5 CTomilPolicy* clone() const

[0156] This method is how a client creates a deep copy of this object. For example: CTomilPolicy* pol2 = pol1->clone();

10 static CTomilPolicy* CreatePolicyFromAB (const char* address, const char* policyLabel)

[0157] This method allows clients to create a policy object using attributes from pre-defined policies contained in a local address book. See CTomilApplication::RestoreDefaultAddressBook() for the default location of the address book. This method will throw a CTomilError if either parameter is bad or no TomIL address book exists. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

15 address: The address token for which this particular policy is defined. For example: "atl".
 policyLabel: The name of the policy token. For example "Hat-email-policy"

20 const int GetTimeOut () const

[0158] This method returns the time-out value for the given policy.

25 const int GetNumRetries () const

[0159] This method returns the number-of-retries value for the given policy.

30 const int GetTBR () const

[0160] This method returns the time-between-retries value for the given policy.

35 const ETomilAckType GetAckType () const

[0161] This method returns the type of acknowledgment specified for the given policy.

TomIL Address Components

[0162] This section contains all the classes necessary to construct CTomilAddress objects. CTomilAddress objects are "generic" from an interface (i.e., external) point of view, but all addresses will contain a protocol-specific implementation (like e-mail, HTTP, FTP, etc.) These protocol-specific implementations are constructed through the factories that are part of this package. Each factory is specific to a protocol and that is where specialization occurs from a client perspective. For example, an e-mail specific factory would take one parameter: an e-mail address, whereas an HTTP-specific factory would take a URL for a parameter and possibly a proxy server address. For the first release of TOM/TomIL there will only be one address protocol: email.

Class --> CTomilAddressFactory

[0163] This is the abstract base class from which all address factories derive. Even though users will not use this class directly, they will use the CreateAddress method via instantiations of the derived factories. Address factories are one of the mechanisms through which clients can create CTomilAddress objects.

Public Methods:

55 virtual CTomilAddress* CreateAddressFromAB(const char* alias) = 0;

[0164] This method allows clients to create well-constructed CTomilAddress objects from parameters gathered in the current TomIL address book under the alias specified by the "alias" parameter passed to this method. This method

will be overridden in every derived factory to return a CTomilAddress object with a <protocol> specific implementation. For the first release of TomIL, the only value of <protocol> is email. This method will throw a CTomilError if the alias parameter is bad or no TomIL address book exists. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

5 alias: a character string that identifies an entry in the current TomIL address book from which the CTomilAddress object will be created. See CTomilApplication::RestoreDefaultAddressBook() for the default location of the address book.

Class --> CTomilAddressFactoryEmail

10 [0165] CTomilAddressFactoryEmail is a specific address factory whose instantiation enables clients to construct well-constructed CTomilAddress objects with implementations that are email specific.

Derived from CTomilAddressFactory

15 **Public Methods:**

static CTomilAddressFactory* Instance()

20 [0166] This is how clients create a CTomilAddressFactoryEmail object. This is the standard protocol for a singleton (127) object. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

Example:

25 [0167]

```
CTomilAddressFactory* af = CTomilAddressfactoryEmail::Instance();
address = af->CreateAddressFromAB( "all" )
CTomilAddress* CreateAddressFromAB( const char* alias )
```

[0168] Implementation of the base class method. See CTomilAddressFactory::CreateAddressFromAB() for details. This method will throw a CTomilError if alias is bad or no TomIL address book exists. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

35 CTomilAddress* CreateEmailAddress(const char* actualAddress)

[0169] This method allows clients to construct well-constructed CTomilAddress object when they have an actual email address. It's important to note that just because a CTomilAddress object is constructed and returned, this makes no claims as to the validity of the email address from which it was constructed. Future releases may include actual run-time validation of email addresses. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

40 actualAddress : A character string representing a valid email address. For example: "rcuser@all.hp.com". Note, it is up to the client to make sure this is a valid email address as TomIL currently has no way to validate this.

45 **Class --> CTomilAddress**

[0170] The CTomilAddress class is a "generic" address-wrapper which contains a protocol-specific implementation set via an address factory.

50 **Public Methods:**

CTomilAddress* clone() const

55 [0171] This method is how a client creates a deep copy of this object. For example:

```
CTomilAddress* add2 = add1->clone();
```

const char* GetAddressString() const

[0172] This method returns a character representation of the address that was used according to the transport protocol chosen. Because protocols vary greatly, they require vastly different addresses to be realized. The string that this method's returned pointer points to should be used for informational/reporting/logging purposes only. There is no format guaranteed on this string, so clients that get in the habit of trying to parse it will be sorely disappointed! The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary

Class --> CTomilAddressStrategy

[0173] CTomilAddressStrategy is the base class from which all address strategy classes derive. Strategy objects in general encapsulate algorithms allowing them to vary independently from the clients that use them. An address strategy is an algorithm that computes the best possible address type for the client to use given the constraints passed to the strategy. For example, a client may want an extremely fast protocol for sending data and may not be concerned so much with reliability. The client would tell a particular address strategy to produce a "FAST" protocol for transporting data and then let the internal algorithms of that strategy compute which protocol would be fastest. Of course, this only makes sense if the strategy algorithm has more than one transport protocol from which to choose. For the first release of TomIL, there is only one protocol, namely email, but for future releases there will be multiple protocols. Introduction of strategy objects at this release is to simply familiarize clients with using them and get the clients used to letting TomIL make strategy decisions for them. The only strategy supported for the first release of TomIL is the "Specific Protocol" strategy described in the next section. It is fully functional.

Public Methods:

CTomilAddressStrategy* clone() const

[0174] This method is how a client creates a deep copy of this object. For example:

```
CTomilAddressStrategy* as2 = as1->clone();
virtual CTomilAddress* CreateAddress( const char* alias )
```

[0175] In each class derived from CTomilAddressStrategy, this method will allow the client to construct a CTomilAddress object from the given alias. This method will throw a CTomilError if alias is bad or no TomIL address book exists. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.
alias: a character string that identifies an entry in the current TomIL address book from which the CTomilAddress object will be created. See CTomilApplication::RestoreDefaultAddressBook() for the default location of the address book.

Class --> CTomilAddressStrategy_SpecificProtocol

[0176] CTomilAddressStrategy_SpecificProtocol is the first and most basic address strategy provided with TomIL. Please read the description of CTomilAddressStrategy before continuing here. A CTomilAddressStrategy_SpecificProtocol strategy is the simplest address strategy shipped with TomIL, and for the first release, the only address strategy. This strategy is constructed with a specific CTomilAddressFactory, like CTomilAddressFactoryEmail for instance. Once constructed, this strategy will return CTomilAddress objects via the CreateAddress() method. The CTomilAddress objects will be constructed according to the constraints of the specific address factory used. If no address factory is passed to the constructor of this strategy, then the default address factory: CTomilAddressFactoryEmail will be used. It may seem silly to have so many supporting objects in factories and strategies when the first release of TomIL only has one protocol: email. But it is imperative to get these objects into the library now so that clients can become familiar with this use model as future releases of TomIL will support multiple protocols.

Derived from CTomilAddressStrategy

Public Methods:

[0177] static CTomilAddressStrategy_SpecificProtocol* Create(
CTomilAddressFactory* factory = 0)

[0178] This is how clients construct a CTomilAddressStrategy_SpecificProtocol object. If the "factory" parameter is not supplied it will default to CTomilAddressFactoryEmail. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

5
CTomilAddress* CreateAddress(const char* alias)

[0179] This method will allow the client to construct a CTomilAddress object from the given alias. The CTomilAddress object will be created according to the factory used to construct this strategy object, or the CTomilAddressFactoryEmail factory if none was specified. This method will throw a CTomilError if alias is bad or no TomIL address book exists. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

10
alias: a character string that identifies an entry in the current TomIL address book from which the CTomilAddress object will be created. See CTomilApplication::RestoreDefaultAddressBook() for the default location of the address book.

15
Example:

[0180]

20
// Create a strategy for getting a CTomilAddress object.
CTomilAddressStrategy_SpecificProtocol* strategy;
strategy = CTomilAddressStrategy_SpecificProtocol::Create(
25
CTomilAddressFactoryEmail::Instance());

CTomilAddress* address;

30
address = strategy->CreateAddress("atl");
// The above line invokes the CreateAddress() method on the "specific protocol"
// strategy. That method will use a CTomilAddressFactoryEmail factory to look
// up the alias "atl" in the current TomIL address book and then retrieve the email

35
// specific address associated with "atl" and construct a CTomilAddress object
// using that address. If any errors are encountered in this process, a CTomilError
// object will be thrown.
40
delete strategy;
delete address;

45 TomIL Job Components

[0181] A Job represents all of the information necessary for a data transmission - what data to send, where to send it and by what protocol, and also all of the transmission policy parameters.

50 Class --> CTomilJob

[0182] The constructor for CTomilJob is protected to force instantiation through the static CTomilJob::Create() method.

55 Public Methods:

[0183] static CTomilJob* Create (const CTomilAddress* address,
const CTomilPolicy* policy)

[0184] This is how clients create a CTomilJob object. If either parameter is NULL, a CTomilError will be thrown. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

5 CTomilJob* clone() const

[0185] This method is how a client creates a deep copy of this object. For example: CTomilJob* job2 = job1->clone();

```
10           enum ETomilDirStructure { TOMIL_FLAT_FILE, TOMIL_FULL_DIRECTORY };
void AddFile (const char* filename, const ETomilDirStructure ds = TOMIL_FULL_DIRECTORY )
```

[0186] The AddFile methods adds a filename to an internal filename list for a given job. The filename will be validated at this point to make sure that it exists, is readable, etc. If the file is invalid, a CTomilError will be thrown. The following 15 are restrictions on the file paths passed to AddFile(). Violation of any of these restrictions will cause a CTomilError object to be thrown.

- Path names to files must be absolute. For HP-UX, this means the pathname to the file must begin with a '/'. UNC names are not allowed in release 1.
- Any file added to a CTomilJob using AddFile() must exist at the time it is added.
- The basename of all files added in conjunction with the TOMIL_FLAT_FILE flag set must be unique. For example, adding /usr/tmp/foo and /etc/foo, each with the TOMIL_FLAT_FILE flag set will cause an error to be thrown. This is because the TOMIL_FLAT_FILE flag causes all files to be laid down on the receiving system using only the file's basename. In this case, /usr/tmp/foo would be delivered as "foo", and then /etc/foo would overwrite "foo" as it was delivered.

ds : A variable that allows the client application to specify how files will be treated on the receiving end when they are put into the receiver's file structure. If TOMIL_FULL_DIRECTORY (the default) is used, files will be put into the receiving 30 file structure according to their full path name (i.e., the file "/usr/lib/foo" would be put into /OUTBOX/<application_name>/<transaction_id>/usr/lib/foo).

[0187] If TOMIL_FLAT_FILE is used, then there is no directory structure constructed on the receiving side (i.e., the file "/usr/lib/foo" would be put into /OUTBOX/<application_name>/<transaction_id>/foo).

[0188] Note: If using this value, clients must NOT construct a job with files which have the same basename, like /usr/lib/foo and /usr/etc/foo. This would cause /usr/lib/foo to be put on the receiver under the name "foo" and then it would 35 be overwritten with /usr/etc/foo. Because this situation causes loss of data, it will be prohibited and a CTomilError will be thrown to the client app if this is detected.

fileName: a character string representing an absolute path to a specific data file. No wildcarding is allowed. No regular expressions are allowed. Note: for NT systems, a volume label must be included (i.e., C:\temp\foo rather than \temp\foo).

40 Class --> CTomilTransactionID

[0189] A CTomilTransactionID object is an object by which any particular job can be identified. Any time a job is transmitted to a destination, a CTomilTransactionID object is returned for that job. (No objects will be returned if TOM 45 re-transmits a job as re-transmission is all under the control of TOM and is completely divorced from TomIL. No matter how many times a job has to be re-transmitted, the initial CTomilTransactionID object will always identify the given transaction.) A CTomilTransactionID object contains a GUID that is guaranteed to be unique across the computing universe.

50 **Public Methods:**

```
static CTomilTransactionID* Create( const char* guid )
```

[0190] Internal users as well as external clients may both want to create a CTomilTransactionID object by passing a 55 GUID string parameter into it. A client may want to do this when bringing a GUID into memory from disk and then using that GUID to construct a CTomilTransactionID object that can be passed to one of the methods of CTomilCourier. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory right before the pointer goes out of scope or a memory leak will be created.

CTomilTransactionID* clone() const

[0191] This method is how a client creates a deep copy of this object. For example: CTomilTransactionID* tid2 = tid1->clone();

5

const CTomilTransactionID& operator= (const char* rhs)

[0192] Assignment operator that lets char*'s be assigned directly to CTomilTransactionID objects. For example:

10 *tid = "guid-identifier-character-string"; // Note the * to de-reference the pointer.

const int operator==(const CTomilTransactionID& rhs)
Standard comparison operator.

15 const int operator!=(const CTomilTransactionID& rhs)
Standard not-comparison operator.

const char* GetID () const

20 [0193] This method simply returns the character string representation of a CTomilTransactionID. This representation will be a GUID. This representation will be "unknown" if the CTomilTransactionID object was constructed with a char* that is 0. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

Class --> CTomilTransactionIDIterator

25

[0194] This class is an iterator that allows clients to get access to all CTomilTransactionID's associated with a particular application interfacing with TOM via TomIL. The only way a valid CTomilTransactionIDIterator object can be created is via a call to CTomilCourier::GetAllIDs() or via a call to clone() on an already valid CTomilTransactionIDIterator object.

30

Public Methods:

[0195] There is no way to create a valid CTomilTransactionIDIterator object except through a call to CTomilCourier:: GetAllIDs() or a call to clone() on an already valid object.

35

CTomilTransactionID Iterator* clone() const

[0196] This method is how a client creates a deep copy of this object. For example:
CTomilTransactionIDIterator* IDIt2 = IDIt1->clone();

40

void Reset() const

[0197] This method will reset an internal pointer to the first TransactionID in the list. This method should always be called before any attempt to traverse a list of TransactionIDs.

45

CTomilTransactionID* Next ()

[0198] This method will advance an internal pointer to the next CTomilTransactionID in the list and return a pointer to that object. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

50

[0199] Code example using a CTomilTransactionIDIterator:

55

```

CTomilCourier* courier = CTomilCourier::Instance();
CTomilTransactionIDIterator* IDIt = 0;
CTomilTransactionStatus* status = 0;
5 CTomilTransactionID* tid;

IDIt = courier->GetAllIDs();
IDIt->Reset();

10 while( tid = IDIt->Next() )
{
    delete status;
    status = courier->Query( tid );
    15 } // for

    delete IDIt;
    delete status;
20

```

Class --> CTomilTransactionStatus

[0200] A CTomilTransactionStatus object contains the status of any particular job that has been transmitted. Once you have a CTomilTransactionStatus object, you can query all the properties of the transaction associated with the CTomilTransactionID object that was passed to CTomilCourier using the myriad of Get*() methods described below. The only way a valid CTomilTransactionStatus object can be created is via a call to CTomilCourier::Query() or via a call to clone() on an already valid CTomilTransactionStatus object.

```

30 enum ETomilTransactionStatusType
{
    35    TOMIL_TRANSACTION_UNKNOWN, // initial state: should NEVER occur.
    TOMIL_TRANSACTION_PENDING,
    TOMIL_TRANSACTION_TIMED_OUT,
    TOMIL_TRANSACTION_COMPLETED,
    TOMIL_TRANSACTION_COMPLETED_WITH_EXCEPTION,
    40    TOMIL_TRANSACTION_SEND_EXCEPTION
}

```

[0201] Those skilled in the art will recognize that an object may be used for this purpose as distinct from a simple enum data type. An enum type is believed to be sufficient in this situation particularly in view of the "optional string" parameter which can be used to exchange information specific to a status.

[0202] A convenience streaming operator is provided for the CTomilTransactionStatus class.
friend class ostream & operator<< (ostream & o, CTomilTransactionStatus*ts);

Public Methods:

[0203] The only way a valid CTomilTransactionStatus object can be created is via a call to CTomilCourier::Query() or via a call to clone() on an already valid CTomilTransactionStatus object.

Example of how to get a CTomilTransactionStatus object

[0204]

```
5      CTomilTransactionStatus* status;
status = Courier->Query( TransactionID );
switch ( status->GetStatus() )
10    {
    case TOMIL_TRANSACTION_PENDING: // do whatever
        break;
    // etc...
    }
15    delete status;
```

const ETomilTransactionStatusType GetStatus () const

20 [0205] This method will return the status that's been assigned to a particular CTomilTransactionStatus object. The status returned will be an enumerated type from the list ETomilTransactionStatusType. For all CTomilTransactionStatus objects that have been created with invalid CTomilTransactionID objects, this method will return the value TOMIL_TRANSACTION_UNKNOWN.

25 const char* GetReturnedString () const

30 [0206] This method will return a pointer to the string that was returned from the process that received the job that was sent. This "returned string" is an optional string that allows receiving applications to communicate other status information back to the sending application. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

[0207] When a CTomilTransactionStatus object is initially created, this string is "" and it will only change if the receiving application had used it to return extra status information to the sending application.

35 const char* GetAppName () const

40 [0208] This method returns the character string representing the application name associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnknownS(). Return values from this method should always be checked against GetUnknownS(). The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

const char* GetAppVer () const

45 [0209] This method returns the character string representing the application version associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnKnownS(). Return values from this method should always be checked against GetUnknownS(). The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

50 const char* GetAppID () const

55 [0210] This method returns the character string representing the application identifier associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnknownS(). Return values from this method should always be checked against GetUnknownS(). The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

const time_t GetTimeSent () const

[0211] This method returns the time_t value (as defined in time.h) representing the time the job was sent which is associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return 0. Return values from this method should always be checked against 0. The return value is specified in UCT: Universal Coordinated Time and should be adjusted by each client for local time.

const ETomilAddressType GetProtocol () const

[0212] This method returns the character enumerated type representing the transport protocol used to send the job associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return TOMIL_ADDRESSSType_UNKNOWN. Return values from this method should always be checked against TOMIL_ADDRESSSType_UNKNOWN. For the first release of TomIL, valid values will always be: TOMIL_EMAIL.

CTomilPolicy* CreatePolicy(void) const

[0213] This method returns a new CTomilPolicy object constructed with the values that were used to create the transaction for which this status object refers. See the section on CTomilPolicy for a description of the methods that can be used to retrieve CTomilPolicy values. Make sure to ALWAYS check this returned pointer for NULL before trying to access any methods of the CTomilPolicy object. This will prevent possible access violations if the policy parameters for this TransactionID were not set for some reason. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

const int GetCurrentTry () const

[0214] This method returns the integer representing the current try (i.e., how many times has this job been sent) value associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnknownI(). Return values from this method should always be checked against GetUnknownI().

const char* GetDestinationAddress () const

[0215] This method returns the character string representing the destination address associated with the given CTomilTransactionID object. There is no format guaranteed on the string returned by this method as it can vary widely between transport protocols. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnknownS(). Return values from this method should always be checked against GetUnknownS(). The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary

const char* GetDestinationAddressAlias () const

[0216] This method returns the character string representing the destination address alias (i.e., The identifier used to retrieve an address from an address book) associated with the given CTomilTransactionID object. If an invalid CTomilTransactionID object was used to create the given CTomilTransactionStatus object, this method will return the same value as the static method: GetUnknownS(). Return values from this method should always be checked against GetUnknownS(). The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

static const char* GetUnknownS ()

[0217] This method returns the character string used for the internal representation of "unknown" for character strings that are returned by the Get*() methods listed above. It is a best practice to always compare any returned string from a Get*() method with this value to make sure the returned value is valid. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

```
static const int GetUnknownValue ()
```

[0218] This method returns the integer used for the internal representation of "unknown" for integers that are returned by the Get*() methods listed above. It is a best practice to always compare any returned integer from a Get*() method with this value to make sure the returned value is valid.

TomIL Courier

[0219] The TomIL Courier is the portal into the transport office. Client applications create and setup jobs and then pass them into the Send() method to be transported via the transport office to their final destination. After sending a package, client applications may asynchronously check the status of a transmission (i.e., if it has been acknowledged yet) by using the Query() method of CTomilCourier.

Class --> CTomilCourier

[0220] The constructor for CTomilCourier is protected to force instantiation through the static CTomilCourier::Instance() method. Think of a CTomilCourier object as the local mailman. A CTomilCourier object handles all communication between an application and TOM. The information that was registered regarding the client application (using CTomilApplication::Register()) will be sent to TOM with every transaction initiated by a CTomilCourier.

[0221] Once the CTomilCourier::Send() method returns, changes made to the job object that it was given will NOT affect that particular transmission. If this behavior is desired, another call to CTomilCourier::Send() must be made.

Public Methods:

```
static CTomilCourier* Instance ()
```

[0222] This method is how a client creates a CTomilCourier object. This is standard protocol for a Singleton(127) object.

For example:

```
CTomilCourier* courier = CTomilCourier::Instance();
```

```
CTomilTransactionID* Send ( const CTomilJob* job, const CTomilApplication* app ) const
```

[0223] This method is the actual communication link between a client and the TOM process running on the client's machine. This method will transmit a CTomilJob to TOM for processing. This method can throw CTomilError objects if communications can not be established with TOM, if either parameter is NULL, or if no files were added to the CTomilJob object passed to this method.

[0224] NOTE: This method (Send()) does NOT block waiting for a job to be sent. It simply streams all data necessary to send a job to TOM and then immediately returns. This method is client-side specific.

[0225] This method returns a new CTomilTransactionID object which is unique to this transaction. The GUID portion of the CTomilTransactionID object must be saved in persistent storage by the client application if it wishes to query the status of the transaction after the CTomilTransactionID goes out of scope. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

```
void Purge ( const CTomilTransactionID* TransactionID ) const
```

[0226] The Purge() method takes a CTomilTransactionID object pointer as a parameter and then removes all references to the CTomilJob object referred to by the given CTomilTransactionID from the TOM persistent and temporary storage. Any future transmissions regarding this CTomilJob are discarded. This method will throw a CTomilError object if communications can not be established with TOM or if the parameter passed in is NULL. This method is client-side specific. This method does NOT affect any processing of the job that may be active on the receiving machine.

CTomilTransactionStatus* Query

```
( const CTomilTransactionID* TransactionID,
```

```
const ETomilQueryType Type =
```

```
TOMIL_NON_DESTRUCTIVE_READ ) const
```

[0227] This method takes a CTomilTransactionID object pointer and a query type as parameters and returns a new CTomilTransactionStatus object which contains status information regarding the CTomilJob referred to by the given

CTomilTransactionID object. The second parameter, Type, specifies what type of read the query is to do. The default, TOMIL_NON_DESTRUCTIVE_READ, forces TOM to not cleanup the transaction after the transaction is complete and the client has been notified through this method, thus transferring the burden of purging the transaction back to the client code. The other option TOMIL_DESTRUCTIVE_READ, lets TOM automatically cleanup a transaction after it has completed and the client has queried it so it knows the transaction is complete. This method will throw a CTomilError object if communications can not be established with TOM or if the transactionID is null. This method is client-side specific. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

```
10 enum ETomilQueryType{
11     TOMIL_NON_DESTRUCTIVE_READ,
12     TOMIL_DESTRUCTIVE_READ }
```

CTomilTransactionIDIterator* GetAllIDs (const CTomilApplication* app) const

20 [0228] This method will return a new CTomilTransactionIDIterator object which allows iteration over all CTomilTransactionID's that are valid for the client's particular application specified by the app parameter. This iterator can be useful for state recovery of an application should it die and then be restarted. The client can get a copy of a CTomilTransactionIDIterator from GetAllIDs() and then re-Query() the status of all its jobs which are still active in the TOM. This method will throw a CTomilError object if communications can not be established with TOM or if the app parameter is NULL.
25 This method is client-side specific. Ownership of the memory attached to the returned pointer will belong to the client. Make sure to delete this memory when the pointer goes out of scope.

void Ack (const CTomilTransactionID* TransactionID, const char* message = 0) const

30 [0229] This method transmits an acknowledgment to the process that originally sent the CTomilJob referred to by the CTomilTransactionID passed in. An optional "status" string may be included with the acknowledgment. This status string will be a 255 (or less) character string. This method will throw a CTomilError object if communications can not be established with TOM or if the transactionID is NULL. This method is server-side specific.

[0230] void PartialAck (const CTomilTransactionID* TransactionID,

35 const char* message = 0) const
[0231] This method provides a means to send progress reports about a job from the server application to the client application. PartialAck() transmits a partial acknowledgment to the process that originally sent the CTomilJob referred to by the CTomilTransactionID passed in. An optional "status" string may be included with the acknowledgment. This status string will be a 255 (or less) character string. This method will throw a CTomilError object if communications can not be established with TOM or if the transactionID is NULL. This method is server-side specific.

40 [0232] A partial acknowledgment differs from a full acknowledgment in that the job on the sending end will not be marked as "complete" upon receipt of a partial acknowledgment.

TomIL Error

45 [0233] This section contains the CTomilError class. CTomilError objects can be thrown throughout TomIL.

Class --> CTomilError

50 [0234] This class represents CTomilError objects. CTomilError objects contain an error code and error text representing internal TomIL errors. These objects can be thrown from methods within the TomIL library of objects. It is necessary that clients catch these objects by pointer.

For example:

```

try {
}
catch( CTomilError* e )
{
    // Process the error
    delete e;
}

```

10

[0235] A convenience streaming operator has been provided for errors.
friend class ostream & operator<< (ostream & o, CTomilError* e);

15 **Public Methods:**

[0236] There is no way to create a CTomilError other than to catch one via a "catch" statement or to use the clone() method on an already created CTomilError object.

20 CTomilError* clone() const

[0237] This method is how a client creates a deep copy of this object. For example: CTomilError* err2 = err1->clone();

const int Errno () const

25

[0238] This method will return the error number of a particular CTomilError object.

const char* ErrString () const

30

[0239] This method will return a pointer to a string that contains the text associated with a particular CTomilError object. The text for all possible CTomilError's will be stored in a TomIL message file. This message file can be localized once it's at a customer's site. The location of all TomIL message files is: /opt/TOM/lib/TomIL.mes on HP-UX, and c:\Tomilib\TomIL.mes on NT. The returned pointer by this method points to memory owned by TomIL. No client cleanup is necessary.

35

[0240] While the invention has been illustrated and described in detail in the drawings and foregoing description, such illustration and description is to be considered as exemplary and not restrictive in character, it being understood that only the preferred embodiment and minor variants thereof have been shown and described and that all changes and modifications which fall within the scope of the appended claims are desired to be protected.

40

Claims

1. A system for managing exchange of data between secured computing systems CHARACTERIZED IN THAT said system comprises:

45

a transmission manager (106, 1204) operable in a first computing system (100, 1200) to manage failure processing in the exchange of said data;

a reception manager (112, 1254) operable in a second computing system (102, 1250) operable in conjunction with said transmission manager to receive said data from said transmission manager and to manage failure processing in the exchange of said data;

a communication medium (150, 200, 1290, 1350..1354) usable by said transmission manager and by said reception manager, wherein said transmission manager utilizes transmission protocols (1300..1308) over said communication medium capable of communicating between said secured computing systems;

a transmit job database (1212) associated with said transmission manager for retaining information regarding an information exchange between said transmission manager and said reception manager; and a receive job database (1260) associated with said reception manager for retaining information regarding an information exchange between said transmission manager and said reception manager.

55

2. The system of claim 1 wherein said transmission manager includes:

a failure monitor (1208) to detect failure in exchange of said data between said transmission manager and said reception manager; and
5 a retransmitter (1210), responsive to said failure monitor, to retransmit said data to said reception manager in response to detection of said failure in exchange of said data.

3. The system of claim 2

10 wherein said transmit job database includes timeout information (1214) indicative of a maximum time for expected completion of said exchange of data, and
wherein said failure monitor includes a timeout detector (1209) to detect failure of said exchange of data by time of said exchange of data exceeding said maximum expected time.

15 4. The system of claim 1 wherein said transmission manager includes:

an acknowledge receptor (1256) for receiving acknowledgment from said reception manager of successful transmission of said exchange of data.

20 5. The system of claim 1 wherein said data transmitted from said transmission manager includes indicia of an originating process of said data, and

25 where said receive database includes:
an application mapping file indicative of an intended destination process (1252) associated with said originating process (1262).

25 6. The system of claim 5 wherein said reception manager includes:

a destination process originator (1258) for starting said destination process in response to receipt of said data.

30 7. A method for exchanging data between secured computing systems CHARACTERIZED IN THAT said method comprises the steps of:

35 receiving a request to transmit data from an originating process in a first system of said secured computing systems (1500);
storing job information in a transmit job database associated with said first system regarding said request to transmit (1502, 1508);
transferring said data from said first system to a second system of said secured computing systems using a communication medium capable of communicating between said secured computing systems (1506); and
assuring successful transfer of said data using said job information in said transmit job database (1600..1622).

40 8. The method of claim 7 wherein the step of assuring successful transfer comprises the step of:

determining that the transfer succeeded in response to receipt within said first system of an acknowledgment from said second system (1620..1622).

45 9. The method of claim 8 wherein said transmit job database includes timeout information indicative of a maximum time for completion of said transfer, and wherein the step of assuring successful transfer further comprises the step of:

50 determining that the transfer failed in response to expiration of said maximum time without receipt of said acknowledgment (1604); and
retransmitting said data in response to a determination that said transfer failed (1606..1608).

10. The method of claim 9 further comprising the step of:

locating information corresponding to said originating process in a receive job database associated with said second system, said information including indicia of a destination process for processing of said data (1716).

55 11. The method of claim 10 further comprising the steps of:

initiating said destination process in response to receipt of said data (1720).




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11

FIG. 12





FIG. 13


FIG. 14

FIG. 15




FIG. 16




FIG. 17





FIG. 18



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 30 4516

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	US 5 222 061 A (DOSHI BHARAT T ET AL) 22 June 1993 (1993-06-22) * abstract * * column 1, line 60 - column 2, line 15 * * column 2, line 43 - column 4, line 39 * * column 5, line 41 - line 51 * + figure 1 * -----	1-11	H04L29/14 H04L1/18
A	EP 0 377 136 A (IBM) 11 July 1990 (1990-07-11) * abstract * * page 2, line 50 - page 3, line 48 * * figures 1,5 *	1,7	
A	US 5 185 795 A (BRIGHT MICHAEL W) 9 February 1993 (1993-02-09) * abstract * * column 2, line 13 - line 25 * * column 5, line 13 - column 6, line 25 * * figures 2,3 *	1,7	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04L
The present search report has been drawn up for all claims			
Page of search EPO FORM 1593/03 02 (P02C01)	Date of completion of the search	Examiner	
THE HAGUE	20 October 1999	Poggio, F	
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	
X : particularly relevant if taken alone			
Y : particularly relevant if combined with another document of the same category			
A : technological background			
O : non-written disclosure			
P : intermediate document			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 4516

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
 The members are as contained in the European Patent Office EDP file on
 The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

20-10-1999

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5222061	A	22-06-1993	NONE	
EP 0377136	A	11-07-1990	US 4970714 A DE 68925958 D DE 68925958 T JP 2228147 A JP 2503086 B	13-11-1990 18-04-1996 14-11-1996 11-09-1990 05-06-1996
US 5185795	A	09-02-1993	NONE	

